Universität Bremen

Fachbereich 3: Mathematik und Informatik

arXiv:1409.2073v1 [cs.CL] 7 Sep 2014

Bachelor Report

An NLP Assistant for Clide

Tobias Kortkamp

Matriculation No. 2491982

Monday 26th May, 2014

First reviewer: Prof. Dr. Rolf Drechsler
Second reviewer: Dr. Berthold Hoffmann
Additional advisors: Dr. Mathias Soeken and Dipl.-Inf. Martin Ring

Selbstständigkeitserklärung

Hiermit erkläre ich, dass ich die vorliegende Arbeit selbstständig angefertigt, nicht anderweitig zu Prüfungszwecken vorgelegt und keine anderen als die angegebenen Hilfsmittel verwendet habe. Sämtliche wissentlich verwendete Textausschnitte, Zitate oder Inhalte anderer Verfasser wurden ausdrücklich als solche gekennzeichnet.

Bremen, den 26.05.2014

Tobias Kortkamp

Contents

1 Introduction **7**

2 Basics **13**
2.1 Clojure . . . 13
2.1.1 Logic programming with core.logic . . . 14
2.1.2 Tawny-OWL . . . 16
2.2 CoreNLP . . . 16

3 Approach **17**
3.1 Architecture . . . 17
3.2 Reconciler . . . 18
3.2.1 Chunking . . . 19
3.2.2 Incorporating text changes . . . 20
3.2.3 Chunk annotations . . . 21
3.3 Integration of CoreNLP . . . 23
3.3.1 Accessing sentences of a text . . . 24
3.3.2 Word maps . . . 24
3.3.3 Coreferences . . . 25
3.4 Building an NLP knowledge base . . . 26
3.5 *Clide* annotations . . . 29
3.5.1 Annotation streams . . . 29
3.5.2 Annotation levels . . . 30
3.5.3 Annotations provided by *clide-nlp* . . . 31

4 Triples **35**
4.1 Triple builders . . . 35
4.2 Reifying triples . . . 44
4.3 Exporting an OWL ontology . . . 49

5 Use case: Graph creation from a natural language specification **51**
5.1 Triple walks . . . 52
5.2 Implementation . . . 55

6 Conclusion **57**

A Part of Speech Tags **59**

B Installation notes **61**

List of Figures **63**

Bibliography **65**

1. Introduction

While developing software or hardware natural language texts are part of the process and used to record or specify the system requirements. The problem is to automatically extract the information contained in these texts and make them available for processing. There are some tools that support developers who want to work with and extract information from these specifications. We are missing an approachable way to define simple rules to automatically extract and use information from a text, that is not written in a severely restricted subset of English.

Some tools like e.g. Cucumber[1], a behavior-driven development framework, solve this by only supporting a DSL.[2] It provides a DSL called Gherkin, which allows users to write test scenarios that both computers and humans can understand. Scenarios consist of steps and each step is parsed using a user provided regular expression [22]. As a consequence a step's regular expression is coupled with the specific phrasing that is used in the step definition. A slight variation in its phrasing without updating the corresponding regular expression or adding a new regular expression, might break the scenario because the provided regular expression does not match the step anymore. Ideally, we would like for two steps, with slightly different phrasing but the same information content, to yield the same output and not break the scenario.

Defining and refining these extraction rules is not a solitary activity, but a collaborative one. By integrating such a system with *Clide*,[3] a development environment with collaboration baked in, we support this aspect from the start.

Clide is a project that was developed as part of Martin Ring's diploma thesis in 2013 at the University of Bremen. It was originally intended to be a web-based development environment for Isabelle[4] only [14, 19]. In contrast with previous Isabelle interfaces, it provides better visualization of the prover's results than traditional and sequential REPL[5]-based interfaces through leveraging Web technologies like HTML5 and JavaScript [19]. It has since undergone further development and has evolved to facilitate collaborative editing of documents with support for other languages besides Isabelle [20].

Clide documents are annotated by assistants. It uses an approach called *Universal Collaboration* where an assistant is seen as just an additional collaborator by the system [20]. While *Clide* is distributive and asynchronous in nature, it provides an interface that can be used to implement assistants "in a simple, synchronous manner" [20].

[1] Available at `http://cukes.info/`
[2] Domain Specific Language
[3] Available at `https://github.com/martinring/clide2`
[4] An interactive theorem prover, available at `http://isabelle.in.tum.de/`
[5] Read Eval Print Loop

A *Clide* assistant is informed about state changes of *Clide*'s environment, be it newly opened documents or edits of a document. It can provide domain specific annotations for any parts of a document's text.

This report describes *clide-nlp*, an NLP[6] assistant for *Clide*. The assistant has the following goals:

- Create a framework for extracting ontologies from a text, by
  - creating an NLP knowledge base (see Chapter 3), and
  - using simple queries on that knowledge base to extract useful information from the text (see Chapter 4).
- Provide annotations for interacting with a text to assist in developing of the queries, including showing
  - the semantic graph of a sentence (see Figure 1.1),
  - the coreferences[7] of the text, and
  - the ontology extracted from the text (see Figure 1.2).
- Work in the collaborative environment that *Clide* provides and keep the ontology and annotations up to date and in sync with text changes.
- Being one of the first *Clide* assistants not developed by *Clide*'s author, *clide-nlp* should also be seen as a test of the limits of *Clide*'s current API for developing assistants. *Clide*'s development continues in parallel with the work on this report.

Chapter 3 introduces the components that constitute *clide-nlp* and how they interact with each other. Section 3.1 and Section 3.5 discuss how *clide-nlp* is integrated into *Clide*.

Because of *Clide*'s interactive and collaborative nature, *clide-nlp* has to contend with continuously changing text. Section 3.2 discusses a model for providing support for incorporating text changes into *clide-nlp*'s internal model. The provided annotations need to reflect these changes immediately, or as fast as possible.

clide-nlp uses the CoreNLP framework by the Stanford University's NLP group that provides access to several NLP tools, including a coreference resolution system, which groups related entities in a text together, and a dependency parser, that describes the relations between sentence parts. Section 3.3 and Section 3.4 describe how we can build an NLP knowledge base based on these tools.

Chapter 4 shows how we can leverage the knowledge base to extract an ontology through a series of simple queries. The queries are described in detail in Section 4.1. Section 4.2 describes how the queries' results are used to create an ontology. Section 4.3 discusses one possible way of how the ontology can be exported to an OWL[8] ontology.

[6] Natural Language Processing
[7] Groups of related entities in a text
[8] Web Ontology Language

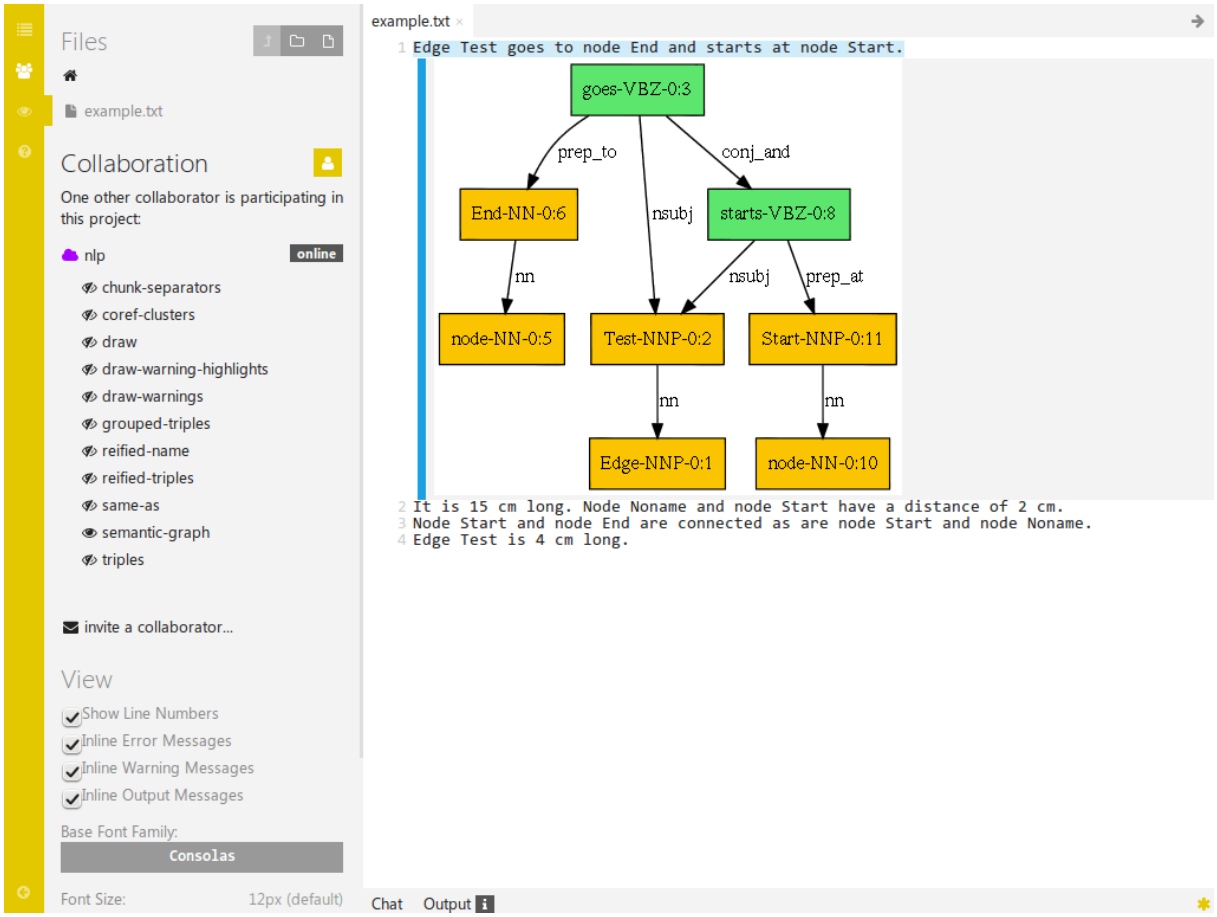

Figure 1.1.: The *clide-nlp* annotation semantic-graph showing the semantic graph for the current sentence (highlighted in blue).

In Chapter 5 we describe an example application that uses the extracted ontology. We extend *clide-nlp* to create graphs from natural language specifications by leveraging the ontology. Figure 1.3 shows an example graph.

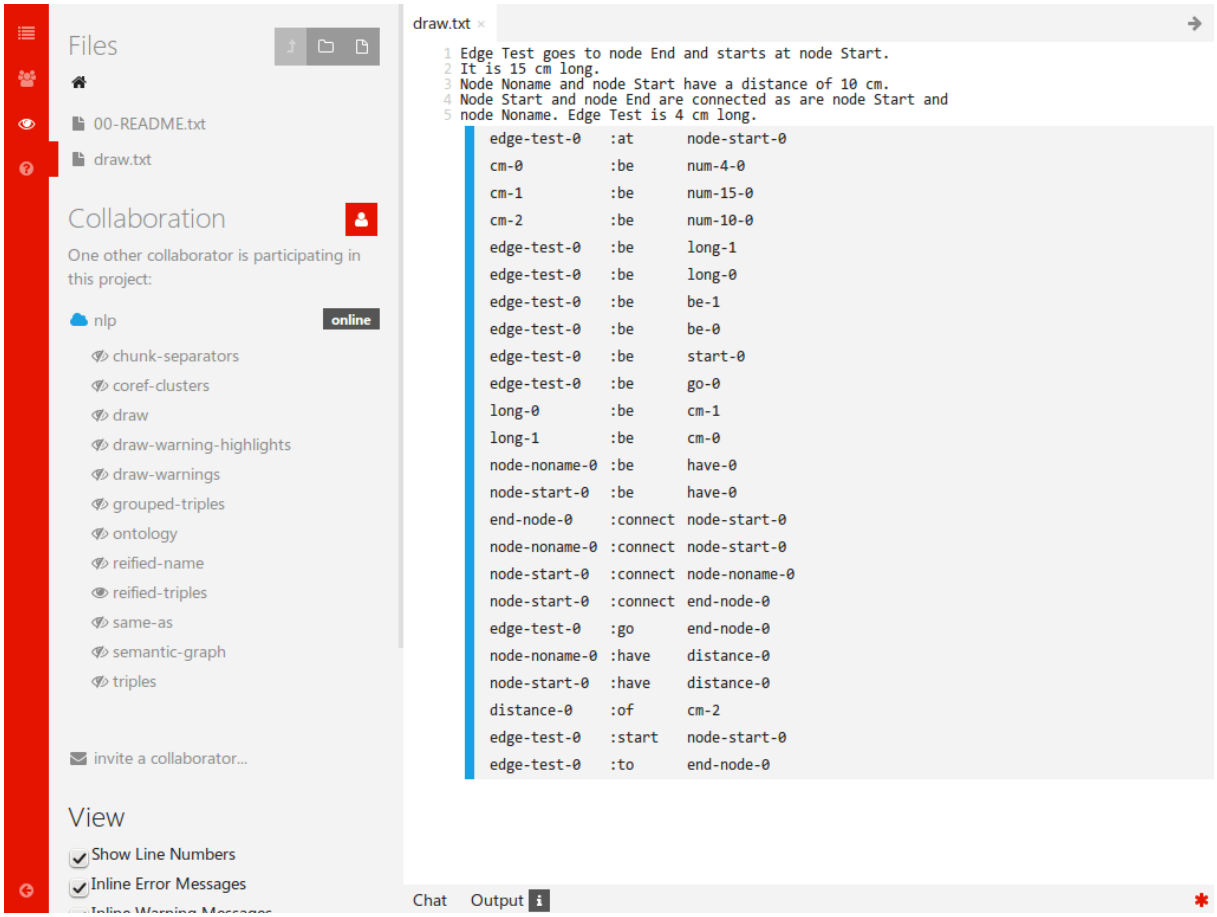

Figure 1.2.: The *clide-nlp* annotation reified-triples showing the triples that are found for the given text and the annotation same-as highlighting all words that are detected as belonging together in red. The highlighted words are part of the group Edge-Test-it-0.

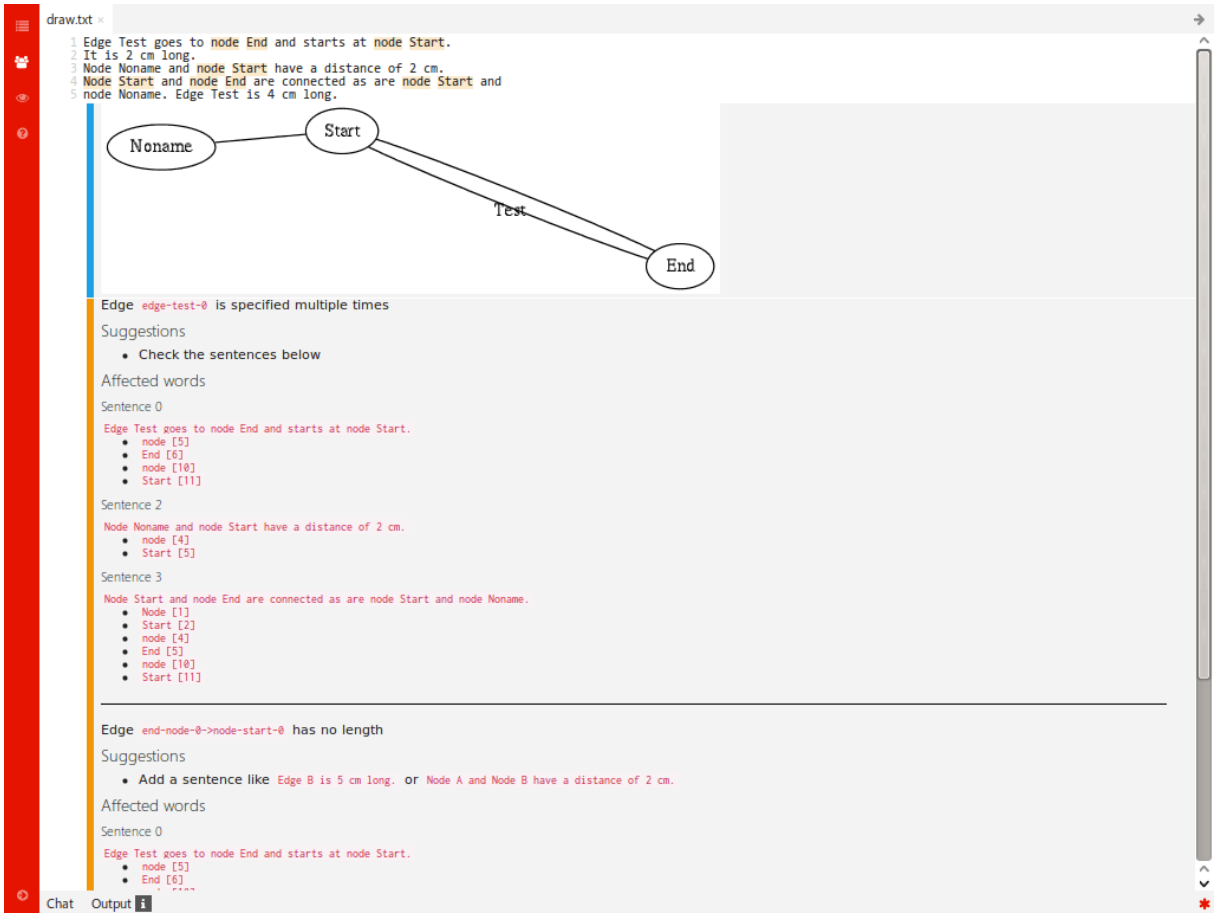
draw.txt
1 Edge Test goes to node End and starts at node Start.
2 It is 2 cm long.
3 Node Noname and node Start have a distance of 2 cm.
4 Node Start and node End are connected as are node Start and
5 node Noname. Edge Test is 4 cm long.
Noname
Start
Test
End
Edge edge-test-0 is specified multiple times
Suggestions
Check the sentences below
Affected words
Sentence 0
Edge Test goes to node End and starts at node Start.
node [5]
End [6]
node [10]
Start [11]
Sentence 2
Node Noname and node Start have a distance of 2 cm.
node [4]
Start [5]
Sentence 3
Node Start and node End are connected as are node Start and node Noname.
Node [1]
Start [2]
node [4]
End [5]
node [10]
Start [11]
Edge end-node-0->node-start-0 has no length
Suggestions
Add a sentence like Edge B is 5 cm long. or Node A and Node B have a distance of 2 cm.
Affected words
Sentence 0
Edge Test goes to node End and starts at node Start.
node [5]
End [6]
Chat
Output

Figure 1.3.: Showing all annotation from the example application draw, draw-warnings and draw-warning-highlights that draws the graph specified in the text and highlights potential problems and ambiguous values in the text.

2. Basics

This chapter introduces some of the concepts needed to understand this report.

clide-nlp is implemented in Clojure. We first give a short overview of Clojure and its logic programming library core.logic. We conclude this chapter with a short introduction of CoreNLP, the NLP framework used by *clide-nlp*.

2.1. Clojure

Clojure is a functional programming language and has several features that make it a good fit for NLP applications.

Clojure provides a REPL which allows us to patch in new code into a running system, enabling a small "thought-code-feedback loop" [7].

It provides built-in persistent data structures with their own reader syntax, that are immutable and have clear and intuitive value equality semantics (1 is 1 and [1 2 3] is [1 2 3]).

Clojure is built on top of the JVM and has good interoperability support with Java libraries. This allows us to leverage all existing Java NLP libraries [6, 21].

Clojure's macros can be used to extend the language when needed. This is used heavily by core.logic, which adds logic programming to Clojure.

The following table summarizes the aspects of Clojure's syntax that is important for reading this report:

Type	**Example**	**Description**
Function definition	(defn f [x y] ...)	defn defines a function. Here we define the function f that takes 2 arguments x and y.
Function call	(f x)	Clojure is a Lisp and uses prefix notation. Here we call the function f with argument x.
Keyword	:a	Keywords are often used as keys in a map, because they evaluate to themselves and can be used as functions that look themselves up in the associative data structures (e.g. a map) that is passed to them as their first argument.
Map	{:a "Hello" :b 0}	A map with the key-value pairs (:a, "Hello") and (:b, 0)
Vector	[1 2 "Hello"]	A vector with the elements 1, 2 and "Hello"

There is no special interoperability support for interfacing with Scala libraries, like e.g. *Clide*.

Interfacing with Scala directly from Java is already challenging, and interfacing with Scala from Clojure adds an additional complication. The work around is to write the components that use *Clide* directly in Scala and use Clojure's runtime interface to call into Clojure code from Scala.

More information on Clojure is available in [6, 7].

2.1.1. Logic programming with core.logic

core.logic[9] adds logic programming capabilities to Clojure. It is based on miniKanren, a logic programming library for Scheme, developed by William E. Byrd as part of his PhD thesis [2].

Because core.logic is a library for Clojure, we can mix functional and logic programming freely, dropping down to core.logic when we need it and use Clojure functional programming aspects otherwise [21].

We summarize the most important functions and macros of core.logic here:

Type	Example	Description
Run a query	(run* [q] ...)	Runs a core.logic query by trying to unify a value with q. Returns a list of all possible values for q
Create logic variables	(fresh [a b] ...)	Creates two unbound logic variables a and b
Unnamed logic variable	(lvar)	Returns a new logic variable
Logical disjunction	(conde [<branch1>] [<branch2>] ...)	Tries all branches consecutively
Soft cut	(conda ...)	Like conde but stops the search as soon as a branch succeeds
Feature extraction	(featurec {:a 4 :b 5} {:a q})	Extracts features from maps. The example binds the logic variable q to 4.
Unify	($\equiv$ q 4)	Unifies the logic variable q with 4.
Never unify	($\not\equiv$ q 4)	Adds a constraint to the logic variable q that it can never be 4.
List membership	(membero q [1 2 3])	A goal that succeeds if q is bound to value that is in the vector [1 2 3]
Extract a logic variable's value inside a query	(project [q] ...)	Extracts the value that is bound to the logic variable q. While in scope of a project q is a regular Clojure value and we can use regular Clojure functions with it.
Domain constraint	(in q (interval 1 10))	Makes sure that q is bound to a value in the interval $[1, 10]$.

[9] Available at `https://github.com/clojure/core.logic`

The presentation of the core.logic code shown here is based on the code presentation in [2, 8].

Presentation	Actual code	
<x>°	`<x>o`	A goal is written with a suffix `o` to distinguish it from already defined functions on the functional programming side, while making clear that they have the same outcome in both paradigms [2]. E.g. cons° in core.logic and cons in Clojure
run*	`run*`	
cond^e	`conde`	The branching macros have an added suffix to distinguish them from the built-in cond.
cond^a	`conda`	
≡	`==`	
≢	`!=`	

The Reasoned Schemer [8] provides a good introduction to miniKanren and in extension also core.logic.[10]

Example. We create a new knowledge base and populate it with three facts about animals.

Relation	Kind	Name
animal	cat	Felix
animal	cat	Mittens
animal	dog	Waldo

Using this knowledge base, we can define a new goal that succeeds iff an animal is a cat or a dog. We make use of the predefined goal member° to check if the value of a logic variable is inside of a collection.

```
1 (defn cat-or-dog° [name q]
2   (fresh [t]
3     (animal t name)
4     (member° t ["cat" "dog"])
5     (≡ t q)))
```

We can then build a query to check what kind of an animal Waldo is:

```
1 (run* [q]
2   (cat-or-dog° "Waldo" q))
3 ⇒ ("dog")
```

Waldo is a dog! And Benjamin?

```
1 (run* [q]
2   (cat-or-dog° "Benjamin" q))
```

[10] Their differences are described on core.logic's Wiki available at `https://github.com/clojure/core.logic/wiki/Differences-from-The-Reasoned-Schemer`.

```
3 ⇒ ()
```

Benjamin does not exist in the knowledge base, so the query returns no result.
We can run the query in reverse to get all cats:

```
1 (run* [q]
2   (cat-or-dogo q "cat"))
3 ⇒ ("Felix" "Mittens")
```

While core.logic supports relational programming, our usage of several non-relation goals, like cond^a or project, makes all of our core.logic usage effectively non-relational.

2.1.2. Tawny-OWL

Tawny-OWL[11] is a Clojure library that provides a domain specific language for building OWL ontologies [13].

clide-nlp uses Tawny-OWL for building OWL ontologies out of its custom ontologies it extracts from texts (see Section 4.3). Exporting OWL ontologies allows us to make use of the existing OWL tools, like e.g. querying them via SparQL [10].

2.2. CoreNLP

CoreNLP is an NLP framework that was created by the NLP group at Stanford University.[12] It includes several components that facilitate developing NLP applications or algorithms. *clide-nlp* uses CoreNLP's dependency parser and its coreference resolution system.

The dependency parser makes the underlying structures of sentences visible in the form of grammatical relations between sentence parts [4]. The output of this component can be modeled as a graph, where the grammatical relations are the graph's edges and the nodes are the sentence parts. We call this graph *semantic graph* in this report. Examples of semantic graphs are available in Section 4.1 and Section 3.3. The grammatical relations are described in [5]. The dependency parser can collapse prepositions and coordinations into grammatical relations [4, 5]. *clide-nlp* uses this feature to simplify the resulting semantic graphs.

The dependency parser makes use of CoreNLP's part-of-speech (POS) tagger. Its tagset is based on the POS tagset used by the Penn Treebank, described in [15]. Because semantic graphs contain POS tags and *clide-nlp* makes heavy use of semantic graphs, Table A.1 provides an overview over some of the tags provided by CoreNLP.

The deterministic coreference resolution system is used to identify entities refer to each other (also called mentions). It was introduced in [11, 12, 17, 18]. It competed in the CoNLL Shared Task 2011, where it achieved the highest score in both the closed and open tracks [12].

Section 3.3.3 goes into more detail and shows an example coreference cluster.

[11] Available at `https://github.com/phillord/tawny-owl`
[12] Available at `http://nlp.stanford.edu/`

3. Approach

This chapter introduces the components that constitute *clide-nlp* and how they interact with each other. We first describe how data flows between the components by giving a high-level architecture overview in Section 3.1. We then delve deeper into the implementation.

clide-nlp receives a continuous stream of events (text operations, cursor movements, and annotation requests) from *Clide* that need to be integrated with *clide-nlp*'s underlying computation model. Section 3.2 describes this model and how we integrate changes into it.

In Section 3.3 we continue with a description of how CoreNLP is integrated into the system by revealing some pitfalls that occur when using the diverse data structures provided by CoreNLP and how we can avoid them. In Section 3.4 we then build a knowledge base from the data provided by CoreNLP which we can use with core.logic and that forms the basis for the remaining chapters.

We conclude this chapter with Section 3.5 and explain the user-facing side of *clide-nlp* by showing what annotations are provided and what caveats apply given the current annotation model of *Clide*.

3.1. Architecture

Figure 3.1 shows how data flows between the components in *clide-nlp*.

Clide assistants need to provide a subclass of `AssistantServer`. `AssistantServer` has support for connecting and receiving messages from *Clide* built-in and abstracts away the underlying Akka implementation.

clide-nlp calls its `AssistantServer` subclass `AsyncAssistantServer`. It receives events for file changes and cursor movements.

All file changes are passed to the reconciler, which uses referentially transparent functions. We pass them the current state of a file and it returns an updated version of that state. The `AsyncAssistantServer` instance is responsible for storing that state and retrieving it when needed.

The `AsyncAssistantServer` enqueues the reconciler state and the cursor position of the change in a queue that is used by the annotation-loop.

The annotation-loop reads one element (state and position) at a time from the queue and prepares *Clide* annotations based on it. See Section 3.5.3 for a list of provided annotations.

The annotations are not computed in the annotation-loop, but are sent to the assistant-loop,

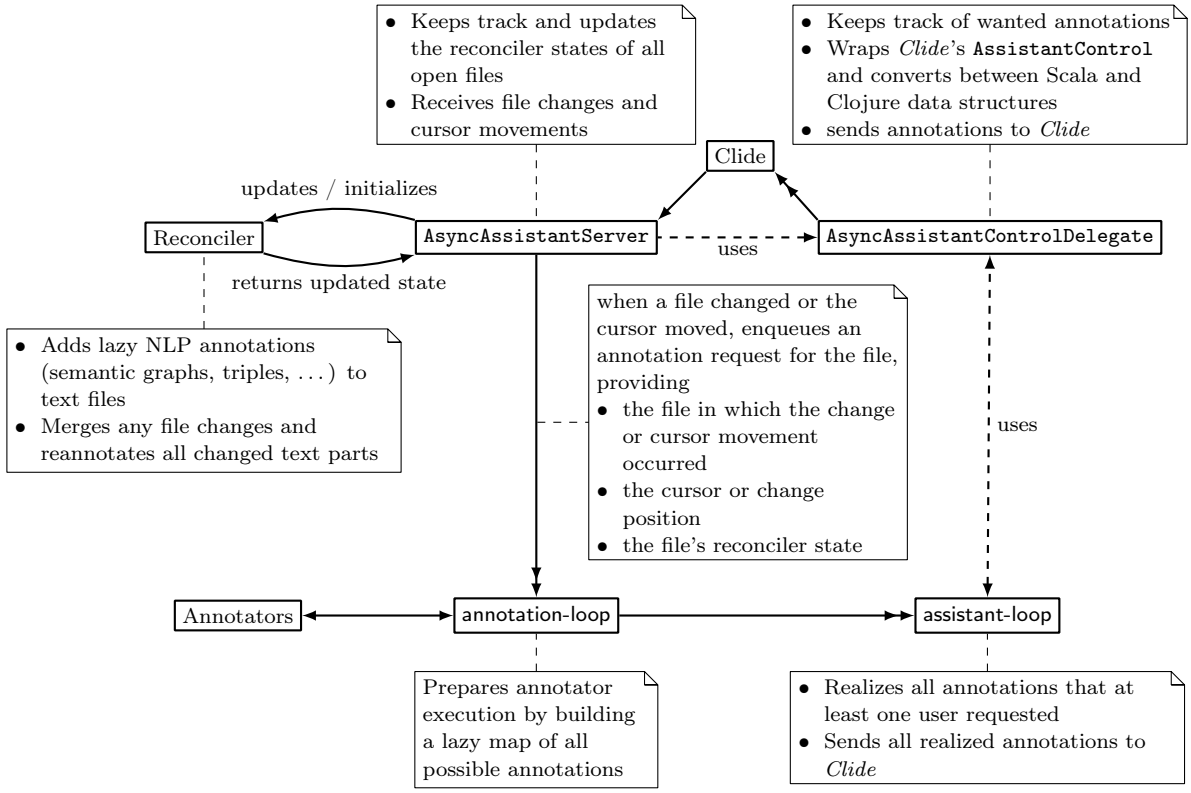
• Keeps track and updates the reconciler states of all open files
• Receives file changes and cursor movements
• Keeps track of wanted annotations
• Wraps Clide's AssistantControl and converts between Scala and Clojure data structures
• sends annotations to Clide
Clide
updates / initializes
Reconciler
AsyncAssistantServer
uses
AsyncAssistantControlDelegate
returns updated state
• Adds lazy NLP annotations (semantic graphs, triples, ...) to text files
• Merges any file changes and reannotates all changed text parts
when a file changed or the cursor moved, enqueues an annotation request for the file, providing
• the file in which the change or cursor movement occurred
• the cursor or change position
• the file's reconciler state
uses
Annotators
annotation-loop
assistant-loop
Prepares annotator execution by building a lazy map of all possible annotations
• Realizes all annotations that at least one user requested
• Sends all realized annotations to Clide
A B Data flows from A to B.
A B Data flows from A to B and from B to A.
A B Data is sent from A and is queued at B (asynchronously).

Figure 3.1.: High-level architecture overview showing the data flow between *clide-nlp*'s components

which is responsible for realizing only those annotations that the users of *clide-nlp* want to see and then sending them to *Clide*.

The annotators that create *Clide* annotations make use of the reconciler's lazy NLP annotations (see Section 3.2). The combination of lazy *Clide* annotations and lazy NLP annotations guarantees that *clide-nlp* only does work when it really has to.

3.2. Reconciler

Clide is a collaborative editing environment. Multiple users may change the current file's text at any time. As such *clide-nlp* needs a way to incorporate those text changes into its own data model.

In traditional IDEs the process which incorporates changes into its data models is called *reconciliation*. *Clide* itself does not provide built-in support for reconciling yet. *clide-nlp* has to provide that support itself.

The reconciler has several related tasks, which are performed in the following order:

1. Split the input text into separate chunks
2. Replay *Clide* deltas to incorporate text changes and mark all chunks that have changed
3. Compute NLP annotations for each changed chunk

3.2.1. Chunking

clide-nlp splits an input text into several parts to keep the time needed for (re-)computing the NLP annotations to a minimum and to make testing easier.

In the implementation provided with this report, *clide-nlp* splits an input text at the string `"\n----\n"`.

Example. The following text

```
The cat eats the mouse.

----

The mouse is dead.
```

is split into 3 chunks:

1. The cat eats the mouse.
2. ----
3. The mouse is dead.

Each chunk has an associated span. A span is the offset interval from the beginning of the text. In the example above, Chunk 1 has a span of $[0, 25)$ and Chunk 2 is called a chunk separator.

Because a chunk is just a slice of an input text, more complicated chunkers are possible, and indeed would be more useful and realistic than the very simplistic chunker currently implemented.

We could e.g. treat comment blocks in a Java file as a chunk and the code in between blocks as chunk separators. CoreNLP might not understand tokens or characters used in Java comments and as a consequence, we would need to remove them first. If we replace them with whitespace before passing the comment to CoreNLP, we make sure that we can map to the original text in an easy way by mirroring the spans inside the original text and inside the replacement text.[13]

[13] CoreNLP does something similar in its `cleanxml` annotator to remove XML tags from an input text before parsing it.

Example. The Java comment

```
/**
 * This is a comment
 */
```

can be written as

$$S = \texttt{"/**}\backslash n\texttt{␣*␣This␣is␣a␣comment}\backslash n\texttt{␣*/"}$$

The substring

$$T = \texttt{"This␣is␣a␣comment"}$$

has the span $[7, 24)$ in S. If we replace all special characters not understood by CoreNLP with spaces, we would get the string

$$S' = \texttt{"␣␣␣}\backslash n\texttt{␣This␣is␣a␣comment}\backslash n\texttt{␣␣␣"}$$

The span of T in S' is still $[7, 24)$.

3.2.2. Incorporating text changes

In *Clide* text changes are described as a list of operations that describe the steps needed to transform an old version of a text into a new version. There are three operations [20]:

- `Retain(`n`)`
- `Insert(`s`)`
- `Delete(`n`)`

Since *Clide* is written in Scala `Retain`, `Insert` and `Delete` are implemented using Scala case classes.[14] While it is possible to work with theses classes in Clojure, it is easier to translate them to use Clojure's data structures instead.

The translation is straightforward. For each operation, replace

`Retain(`n`)`	with	[:retain n]
`Insert(`s`)`	with	[:insert s]
`Delete(`n`)`	with	[:delete n]

Example. The operations `[Retain(5), Insert("hallo"), Retain(15)]` are translated into the Clojure data [[:retain 5] [:insert "hallo"] [:retain 15]].

To apply the operations we need to maintain a cursor position starting at 0. The cursor's position is the position in the original text, not the edited text. The original text is immutable and is not changed. The edited text is built by applying each operation sequentially:

- [:retain n] moves the cursor n characters ahead and inserts them in the edited text.

[14] Case classes are algebraic data type constructors and allow pattern matching.

- [:insert s] inserts the string s at the current cursor position in the edited text. This will not move the cursor, because it is relative to the original text.
- [:delete n] deletes the next n characters at the current cursor position and moves the cursor n characters ahead.

Example. Given the input text "This␣is␣a␣test." and the operations [[:retain 9] [:insert "the"] [:delete 1] [:retain 6]], we get the text "This␣is␣the␣test." after applying them.

As the reconciler splits a text into several chunks there are some more concerns to address.

The changes to a text can be

1. inside a chunk separator, i.e. between two chunks A and B[15] in which case there are two possible scenarios:
   a. A remains unchanged and the spans of B and of all chunks that follow it have to be updated.
   b. A and B have to be merged together because the chunk separator is not valid anymore. This would also change the spans and indices of all chunks that follow A and B.
2. inside one chunk C, in which case C's span (end offset) and the span of all chunks that follow C must be updated to reflect the changes.

In the actual reconciler implementation changes between two chunks are detected after applying the edit operations and rechunking the text. If the number of chunks changed, the reconciler is simply reinitialized. This greatly simplifies the implementation, but all NLP annotations are lost and need to be rebuilt. An ideal implementation would have to follow all scenarios above.

3.2.3. Chunk annotations

Each chunk has associated annotations that are updated when a chunk changes. The annotations are added to a Clojure map with lazy evaluation semantics. This model allows the reconciler to remain fast even when there are continuous changes. The chunk annotations are only realized, and thus computed, outside of the reconciler.

The graph in Figure 3.2 shows how the annotations depend on each other. Note that graph is the central nexus of *clide-nlp* that pulls all of its parts together. As such every aspect of *clide-nlp* is mirrored in it.

The annotations have the following meaning:

:text, :corenlp-pipeline The graph's inputs are a previously constructed CoreNLP pipeline[16] and the chunk's text.

:corenlp-annotation A CoreNLP annotation is created based on the input text. The annotation provides access to all primitive NLP constructs we need (which includes the coreference chain, the semantic graphs, all split sentences, and information about each token).

[15] Chunk B follows chunk A.
[16] The pipeline needs to be setup to use CoreNLP's depedency parser and its coreference resolution system.

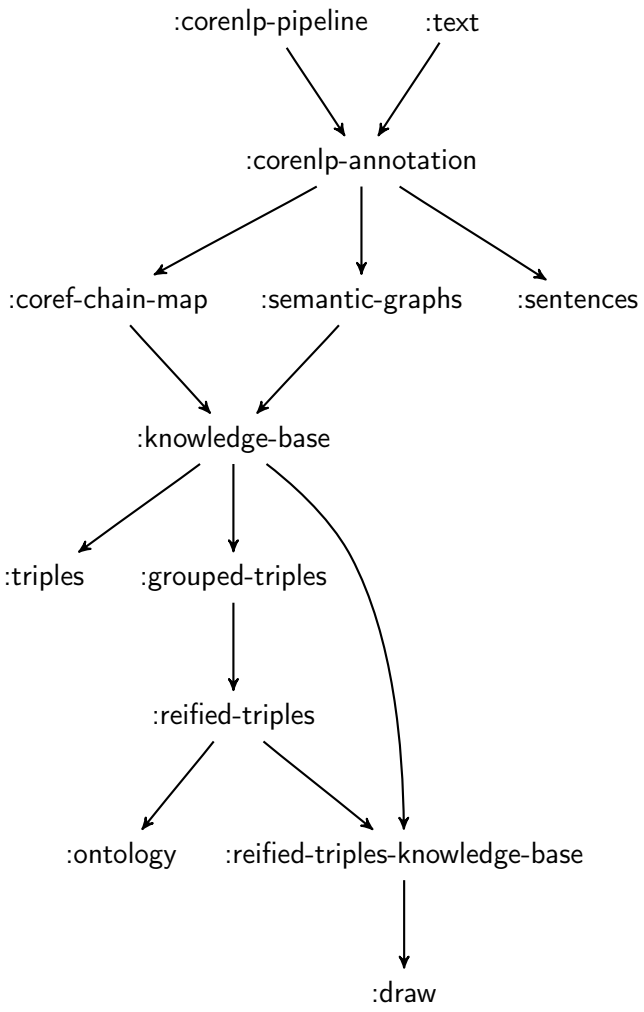

Figure 3.2.: Chunk annotation dependency graph

:semantic-graphs Extracts all semantic graphs from the annotation and creates a Clojure representations of them. In CoreNLP the class that represents nodes in a semantic graph is called `IndexedWord`. All `IndexedWord` instances are mapped to a word map for later use (see Section 3.3.2).

:coref-chain-map Extracts all coreference clusters from the CoreNLP annotation (see Section 3.3.3).

:sentences Extracts information about each of the input text's sentences from the annotation and builds a list of sentence maps (see Section 3.3.1).

:knowledge-base Builds a knowledge base for use with core.logic out of the coref chain map and semantic graphs. The process is described in Section 3.4.

:triples Runs triple builders on the knowledge base that extract useful information gathered from the semantic graphs and builds a list of triples. The triple's subject, predicate, and object only refers to a node in the semantic graph of one sentence. See Chapter 4 for more details.

:grouped-triples Runs triple builders on the knowledge base and groups the triples' subjects

and objects by their coreference cluster. A group is a list of coreferent or otherwise related words. In contrast with the triples extracted by :triples, the grouped triple's subject and object are a list of related words that can span multiple sentences (the whole text) instead of only one sentence. See Section 4.2 for more details.

:reified-triples Multiple grouped triples can all have the same subject or object groups. :reified-triples assigns a unique name to each group i.e. the groups are made real (reified) by giving them a name (see Section 4.2).

:reified-triples-knowledge-base Adds the reified triples to the knowledge base for use by applications that don't want to search the triples sequentially.

:draw An annotation that is introduced in detail in Chapter 5. It tries to draw a graph specified in the input text and to warn about simple ambiguous sentences.

:ontology Builds an OWL ontology (see Section 4.3).

An annotation is only realized when it is directly needed or used by a dependent annotation. As a result, *clide-nlp* does not always have to compute all annotations if the text changed and only does work if it is really needed.

Example. If a client accesses the :knowledge-base annotation only the following annotations are realized: :semantic-graphs, :coref-chain-map, :corenlp-annotation, :text.[17]

3.3. Integration of CoreNLP

clide-nlp tries to not rely on CoreNLP's data structures, because the data structures need to participate in core.logic unification. Due to the mutable nature of CoreNLP's data structures extending them to reliably support unification is problematic.

Additionally, there are inconsistencies in the usages of 0- or 1-based indices in CoreNLP's data structures. This is corrected when constructing *clide-nlp* data structures and allows for easier matching up of the different data structures based on sentence and token indices.

All data structures are implemented using Clojure records. Records are reified maps, which compile down to Java classes. They implement the correct interfaces, so that they can be treated as maps, which we will do from this point on.

We introduce each record by

- giving a short description of its use in *clide-nlp*,
- by listing its available keys with a description of the content of their values,
- and by showing examples with actual data.

[17] Note that :text is the input to the graph and as such always realized.

3.3.1. Accessing sentences of a text

While accessing individual sentences of a text is not important for the core task of *clide-nlp* (extracting an ontology from a text), we need them to create *Clide* annotations that refer to a whole sentence (see Section 3.5).

Sentence maps have the following keys:

:index	The index of the sentence starting at 0.
:span	The 0-based character index span $[a, b)$ of the sentence. The sentence starts at offset a and goes up to offset b.
:text	The text of the sentence.

Example. The text "Felix is a cat. Waldo is a dog. Tweety is a bird." results in the following sentence maps:[18]

:index	0
:span	$[0, 15)$
:text	Felix is a cat.

:index	1
:span	$[16, 31)$
:text	Waldo is a dog.

:index	2
:span	$[32, 49)$
:text	Tweety is a bird.

3.3.2. Word maps

Semantic graph nodes are an integral part of the triple builders introduced in Chapter 4.

The semantic graph node class in CoreNLP is called `IndexedWord`. The information that an `IndexedWord` object provides, is used to create a word map with the following keys:

:sentence	The sentence index this word map refers to. This matches the sentence maps' :index value.
:index	The index of the word's token starting at 1. CoreNLP consistently starts at 1 when counting tokens.
:span	The 0-based character index span $[a, b)$ of the word map.
:tag	The word's part-of-speech tag.
:lemma	The word's lemma.
:token	The word's token.

[18] When checking the spans do not forget to include the spaces between the sentences!

Example. The semantic graph for the input text "Felix is a cat." has the following word maps:

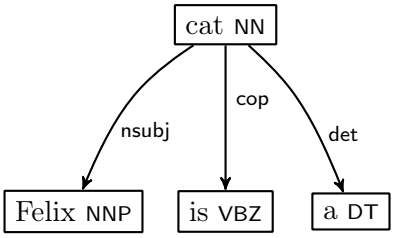

:sentence	0
:index	1
:span	$[0, 5)$
:tag	NNP
:lemma	Felix
:token	Felix

:sentence	0
:index	2
:span	$[6, 8)$
:tag	VBZ
:lemma	be
:token	is

:sentence	0
:index	3
:span	$[9, 10)$
:tag	DT
:lemma	a
:token	a

:sentence	0
:index	4
:span	$[11, 14)$
:tag	NN
:lemma	cat
:token	cat

3.3.3. Coreferences

clide-nlp uses CoreNLP's coreference resolution system to identify which entities in a text are similar to other entities in a text. Coreferences are grouped in clusters. A cluster is made up of mentions and we map them to mention maps with the following keys:

:cluster-id	The coreference cluster id the mention map is a part of.
:sentence	The index of the sentence that contains this mention map. The indices in CoreNLP start at 1 here. We correct them to be 0-based, and as a consequence match a sentence map's :index and a word map's :sentence value.
:index-span	The index span $[a, b)$ refers to the tokens starting at the word map with index a and ends before the word map with index b. The indices are 1-based again, but we do not need to adjust them here, because the token indices of word maps are 1-based, too.
:text	A clear text representation of the mention map.

CoreNLP's coreference system provides additional information about each mention. Information about a mention's gender, its animacy or its number (plural or singular), is however currently unused by *clide-nlp*.

Example. The text "Felix is a cat." has the following coreference cluster and associated mention maps:

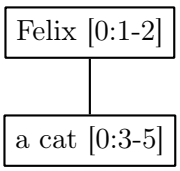

:cluster-id	1
:sentence	0
:index-span	$[1, 2)$
:text	Felix

:cluster-id	1
:sentence	0
:index-span	$[3, 5)$
:text	a cat

We see that the coreference resolution system in CoreNLP identified Felix to have the same meaning as a cat.

3.4. Building an NLP knowledge base

This section describes how the data structures introduced in Section 3.3 are inserted into a core.logic database.

Chapter 4 makes extensive use of the knowledge base and provides usage examples.

We first need to define the relations that we want to provide. They are described in further detail later. *clide-nlp*'s knowledge base provides the following relations:

(word-map w)	provides access to word maps (semantic graph nodes).
(depends dep $reln$ gov)	provides access to semantic graph edges.
(same-as w_1 w_2)	succeeds iff the word maps w_1 and w_2 can be treated as referring to the same word.

Next we need to insert facts into the knowledge base. Given an input text,

1. for every word map w of a semantic graph of a sentence of the text, insert the fact (word-map w).
2. for every semantic graph edge $e = (dep, reln, gov)$ of a semantic graph of a sentence of the text, insert the fact (depends dep $reln$ gov).
3. for the word maps w_1 and w_2 and using word-map, depends and coreference cluster mentions, determine if w_1 and w_2 can be treated as referring to the same word, then insert the facts (same-as w_1 w_2) and (same-as w_2 w_1).

(word-map w)

word-map provides access to the word map w. It is a unary relation. To unify with information from a word map, it needs to be used in conjunction with core.logic's feature extraction goal featurec (see Section 2.1.1).

Example. Using core.logic's featurec goal, we can limit a query to only succeed with maps with specific features. The following query returns all word maps which have a tag NN and an index 0.

```
1 (run* [q]
2   (word-map q)
3   (featurec q {:tag "NN" :index 0}))
```

Because matching a specific set of part of speech tags is used heavily by the triple builders introduced in Chapter 4, the following helper goals are defined:

(tag° w $tags$)	succeeds iff w has one of the tags in the vector $tags$.
(verb° w)	succeeds iff w is a verb, i.e. if it has one of the tags VB, VBD, VBG, VBN, VBP or VBZ.
(noun° w)	succeeds iff w is a noun or pronoun, i.e. if it has one of the tags NNP, NN, NNS. PRP or PRP$.
(wh-word° w)	succeeds iff w is a *wh*-word[19] i.e. if it has one of the tags WDT, WP, WP$ or WRB.

tag° is the basis of the implementations of all of these goals. tag° can be defined in the following way:

```
1 (defn tag°
2   [w tags]
3   (fresh [tag]
4     (word-map w)
5     (featurec w {:tag  tag})
6     (member° tag tags)))
```

(depends *dependent relation governor*)

While word-map provides access to semantic graph nodes, depends provides access to semantic graph edges.

An edge goes from the word map *governor* to the word map *dependent* with the grammatical relation *relation*. We insert every edge of every semantic graph of every sentence of a text into the knowledge base.

relation can be either a string containing a typed dependency relation (see [5]) or if the relation is a collapsed relation, a vector of the first and second part of the relation (e.g. if the relation in the semantic graph was prep_of, it gets split into the vector ["prep", "of"]). Splitting a collapsed dependency in that way keeps core.logic queries simple by making use of its native support for unifying vectors. This allows searching for specific prepositions or all prepositions in a simple manner.

Example. The following query returns the governor and dependent word maps of all edges with a prepositional relation (e.g. prep_of or prep_in) from every semantic graph of a text.

```
1 (run* [q]
2   (fresh [dep gov p]
3     (depends dep ["prep" p] gov)
4     (≡ q [dep gov])))
```

(same-as w_1 w_2)

same-as asserts that the word map w_1 is the same as the word map w_2 and that we can treat

the words as being instances of the same word group.

same-as should be commutative, so that the order of w_1 or w_2 does not matter. If (same-as w_1 w_2) succeeds, (same-as w_2 w_1) succeeds, too.

Because w_1 and w_2 are instances of the same word group and we like the word groups to be about a concrete *thing*, we want to limited them to only include pronouns, nouns, or determiners. Including adjectives e.g. does not make sense because they are properties of word groups. Verbs are used between two or more word groups.

same-as is important for grouping triples (see Section 4.2). There are several aspects for when we can consider two word maps the same.

To be included in the same-as relation, the word maps w_1 and w_2 need to fulfill at least one of the following rules:

- w_1 and w_2 need to map to the same coreference cluster. We need to find the corresponding word maps of each of the cluster's mentions. We can do this in a core.logic query by constraining a word maps index to be inside of the mention's index span:

```
1 (let [[start end] (:index-span mention), sentence (:sentence mention)]
2   (run* [q]
3     (fresh [index tag]
4       (word-map q)
5       (featurec q {:index  index, :tag tag, :sentence sentence})
6       (in index (interval start (dec end))))))
```

  We further limit the query result to only include word maps that represent pronouns, nouns, or determiners.

  We run the query for every cluster mention and select every 2 combination of the found word maps and record the facts:[20]

  (same-as w_1 w_2) (same-as w_2 w_1)
  (same-as w_1 w_1) (same-as w_2 w_2)

  We repeat this process for every coreference cluster.

- The query

```
1 (run* [w1 w2]
2   (noun° w1)
3   (noun° w2)
4   (depends w1 "nn" w2))
```

  succeeds for w_1 and w_2.

  nn is the *noun compound modifier* dependency relation that asserts that one noun modifies another noun [5].

[20] While facts might be recorded twice, we can safely ignore this.

By including compound nouns in same-as, we ensure that every word of a compound noun is assigned to the same word group later (see Section 4.2).

- w_1 and w_2 need to be linked by a *wh*-word. Word maps that are linked by a *wh*-word can be found with the following query:

```
1 (run* [w1 w2]
2   (fresh [w]
3     (wh-word° w1)
4     (depends w1 "nsubj" w)
5     (depends w (lvar) w2)))
```

CoreNLP's coreference system does not include *wh*-words in its mentions. Some triples found by the triple builders in Chapter 4 have a *wh*-word as their subject and we need to make sure that they can be grouped together with the other groups and do not create a group by themselves.

3.5. Clide annotations

Clide annotations are used to provide rich information about specific parts of a text. They are currently static and non-interactive[21]. There is e.g. no way to jump to a specific word in the document from an annotation. While this limits their usefulness, it does not prevent them from being helpful.

They follow the same model as the operations sent by *Clide* (see Section 3.2.2) and are represented as a list of annotation operations, called an annotation stream. In *Clide* they are simply called annotations, but we use the term *annotation stream* to distinguish them from the annotation lists they contain.

3.5.1. Annotation streams

An annotation stream is a list of annotation operations. There are two types of operations:

- `Plain(`n`)`
- `Annotate(`n`, `*annotations*`)`

where n is the annotation's length and *annotations* is a list of tuples $(type, content)$ with *type* being the annotation type and *content* a string containing the actual annotation content.

We again translate the Scala syntax into Clojure data and replace

`Plain(`n`)`	with	[:plain n]
`Annotate(`n`, `*annotations*`)`	with	[:annotate n *annotations*]

To apply an annotation stream we need to maintain a cursor position starting at 0. The annotation is applied by applying each action sequentially and mutating the cursor position afterwards.

[21] This means that you cannot interact with the annotation itself, because they are view-only. You can however request a new annotation.

- [:plain n] skips n characters from the current cursor position c_i, and adds n to it: $c_{i+1} = c_i + n$.
- [:annotate n *annotations*] moves the cursor n characters ahead and annotates the text span from $[c_i, c_{i+1})$ where $c_{i+1} = c_i + n$. *Clide* then interprets the annotation list *annotations* and displays them.

where

i	the index of the operation in the list
c_0	0
c_i	the cursor position before applying the operation
c_{i+1}	the cursor position after applying the operation

There is a direct correspondence between annotation and edit operations (see Section 3.2). Ignoring *annotations*, we can treat :plain and :annotate as `:retain` operations [20].

An annotation stream should span the whole text, that is by summing up the n-s of each operation, we would get the text length.

Clide defines several annotation types. *clide-nlp* uses the following subset of them:

Class sets the CSS class to use for the annotation.

Tooltip sets the tooltip used when hovering over the annotation.

WarningMessage display a warning inline.

Output for displaying generated information, that is not displayed inline by default.

Clide allows HTML inside of its annotations, which means that we can display richer annotations than only simple plain text annotations.

Example. Given the text "The␣cat␣is␣hungry." of length 18 and the annotation stream
[[:plain 4], [:annotate 3 [[:Class "error"]]], [:plain 11]]
and assuming [:Class "error"] is meant to color the annotated text red, applying the annotation stream to the text results in
"The␣cat␣is␣hungry."

3.5.2. Annotation levels

While *Clide* annotations can annotate arbitrary ranges of a text, it is useful to distinguish between different annotation levels in *clide-nlp*:

Text Annotates the whole text

Chunk Annotates a chunk (as defined in Section 3.2)

Word Annotates a single word

Sentence Annotates a whole sentence

Every annotation in *clide-nlp*, with the exception of a text level annotation, is relative to a chunk. By doing so, we keep the annotation creation as simple as possible.

As discussed in Section 3.2, a chunk has an associated span that indicates the chunk's text position inside of the global text. A chunk's internal span begins at 0. CoreNLP never sees the whole text at once, but instead only sees the texts of every chunk separately, so all spans returned by CoreNLP also begin at 0.[22] It follows that the NLP knowledge base is chunk local, too. As the annotations created by *clide-nlp* all use CoreNLP annotations or the NLP knowledge base, we ideally should use chunk local offsets when creating annotation streams. We later project their chunk local offsets to offsets that *Clide* can interpret correctly.

3.5.3. Annotations provided by clide-nlp

In this section we describe each annotation that *clide-nlp* provides by showing an example of how they appear in *Clide*. The annotations make use of one more of the chunk annotations as described in Section 3.2. There might be some overlap in their names, but they should be treated as separate entities. The names are presented to the user by *Clide*, who can enable them individually (see the names to the left with the "eyes" in Figure 1.1).

chunk-separators — Level: Text

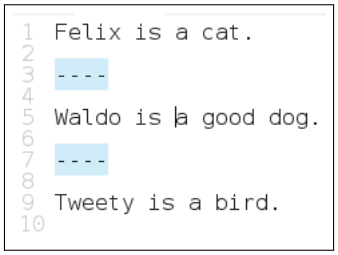

Highlights the chunk separators. Chunks are described in Section 3.2. This makes the chunkers decisions visible.

semantic-graph — Level: Sentence

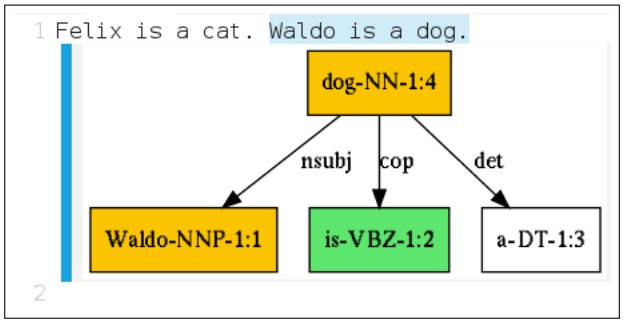

Highlights a sentence and displays the sentence's semantic graph. This annotation is highly useful when creating a new triple builder.

[22] After some correction (see Section 3.3)

coref-cluster — Level: Chunk

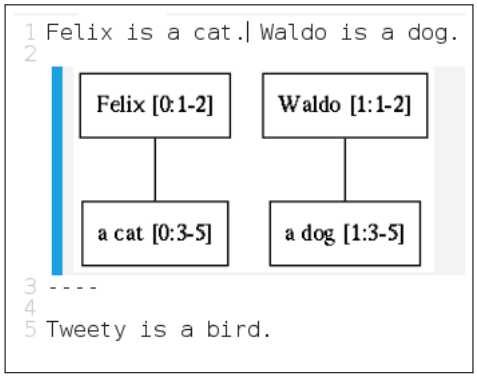

Shows the coref clusters that CoreNLP found for the selected chunk's text.

same-as — Level: Word

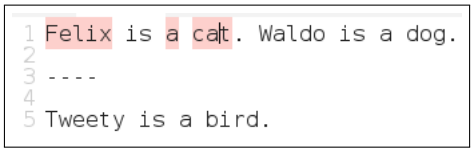

A direct interface to the knowledge base that uses the same-as relation to highlight the related words of the selected word.

triples — Level: Chunk

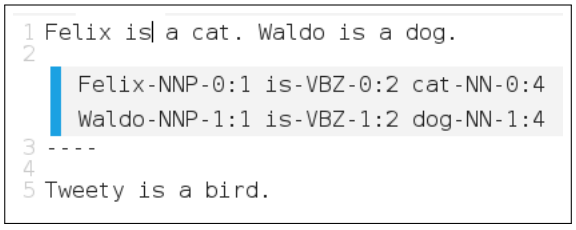

Shows the raw triples extracted from running the triple builders on the selected chunk's text.

grouped-triples — Level: Chunk

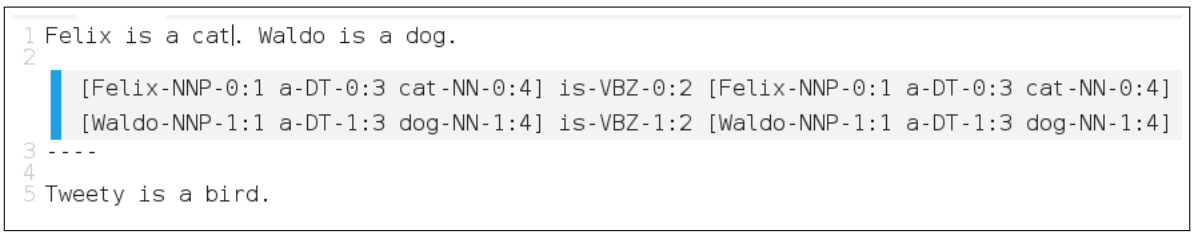
Felix is a cat. Waldo is a dog.
[Felix-NNP-0:1 a-DT-0:3 cat-NN-0:4] is-VBZ-0:2 [Felix-NNP-0:1 a-DT-0:3 cat-NN-0:4]
[Waldo-NNP-1:1 a-DT-1:3 dog-NN-1:4] is-VBZ-1:2 [Waldo-NNP-1:1 a-DT-1:3 dog-NN-1:4]

Tweety is a bird.

Shows the triples' groups.

reified-triples — Level: Chunk

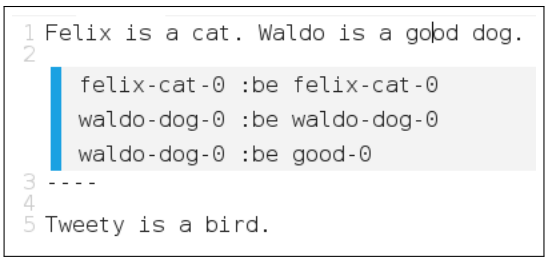

Shows a table with all reified triples extracted from the selected chunk's text.

reified-name — Level: Word

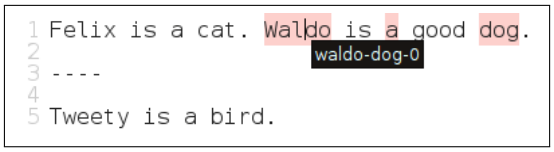

This annotation is identical to the same-as annotation, but additionally adds a tooltip to the highlighted words, showing the word group they belong to.

4. Triples

In the previous chapter we built an NLP knowledge base. This chapter introduces triple builders that runs simple core.logic queries on it and extracts triples. Triples have the form (*subject*, *predicate*, *object*). They represent a unit of useful information.

Where possible a triple's subject and object each have *one* corresponding word map in a semantic graph. Using the knowledge base introduced in the previous chapter, we group the triple's subject and object with other related subjects or objects of other triples, yielding an ontology.

4.1. Triple builders

Triple builders extract triples. Triples have the form (*subject*, *predicate*, *object*). They represent a unit of useful information that is directly extracted from the semantic graph.

A subject usually does something (predicate) with an object. Every subject and object is directly linked to a word map. For predicates however this is not always possible, because there are predicates which are implied by the semantic graph (see the triple builders nsubj-amod or possessive for an example of this). These predicates are called *derived* and written with a colon prefix (e.g. :be) to distinguish them from word map predicates.

While the triple builders only extracts information from semantic graphs, they could be extended to include information from ontologies like WordNet [16], VerbOcean [3] or DBpedia [1] to further constrain the triples that are found.

A triple builder is implemented as a core.logic goal.[23] It either succeeds with a triple or it does not. All triple builders are tied together in a logical disjunction (cond^e) and are all tried in turn. Most of the triple builder's names are directly derived from the semantic graph relation's names they use.

We show each triple builder with

- the actual core.logic query,
- an example input sentence,
- the corresponding semantic graph for the input sentence, with highlights for the relevant nodes and edges to make it easier to follow along,
- a table of the triples that it found when run on the input sentence, displaying the lemmas of its subject, predicate or object,
- and a discussion of the triple builder.

[23] Note that we deviate from the core.logic convention of marking a goal with a postfix ^{o} here.

Note that a triple builder may succeed multiple times and thus may find more than one triple.

nsubj-amod

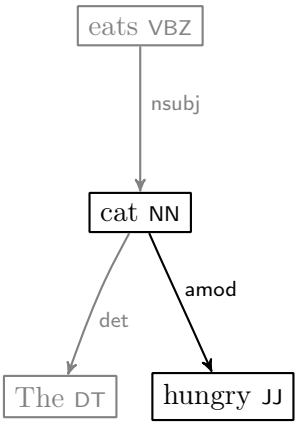

Input text The hungry cat eats.

Query

```
(fresh [subj adj]
  (noun° subj)
  (depends adj "amod" subj)
  (≡ triple [subj :be adj]))
```

Triples found

cat	:be	hungry

The meaning of a subject can be modified with an adjectival modifier (amod) [5].
This triple builder captures the fact that *cat* refers not just to cat but to a *hungry* cat. Because there is no node in the graph that can take the role of a predicate, we introduce a derived predicate :be instead.

nsubj-pred-dobj

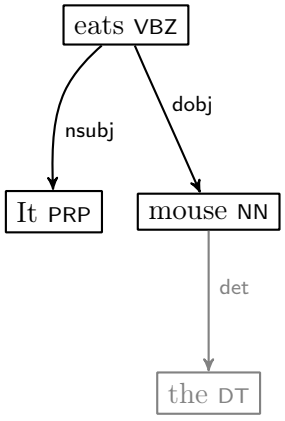

Input text It eats the mouse.

Query

```
(fresh [subject activity object]
  (depends subject "nsubj" activity)
  (depends object "dobj" activity)
  (cond^e
   [(≡ triple [subject activity object])]
   [(fresh [subject2]
      (depends subject2 ["conj" "and"] subject)
      (≡ triple [subject2 activity object]))]))
```

Triples found

it	eat	mouse

This triple builder captures facts from one of the simplest kind of sentences. A subject directly connected to an object with a predicate.

nsubj-VB

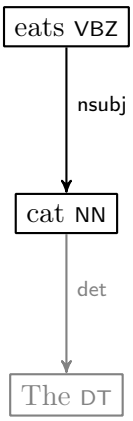

Input text The cat eats.

Query

```
(fresh [subj vb]
  (noun° subj)
  (tag° vb ["VBZ" "VBD" "VBP"])
  (depends subj "nsubj" vb)
  (≡ triple [subj :be vb]))
```

Triples found

cat	:be	eat

Some sentences have intransitive verbs[24] and no object. However, they might still contain useful information.

Here we capture the fact that the *cat eats*. To do this, the verb becomes our triple object and we use a derived predicate :be. This mirrors the behavior used in e.g. the triple builder nsubj-amod for adjectives.

We limit ourselves to sentence with verbs that are in the past tense (tag VBD) or singular present (tags VBP and VBZ) to not always trigger this triple builder for all sentences with an nsubj relation between a noun and a verb.

nsubj-adj-cop

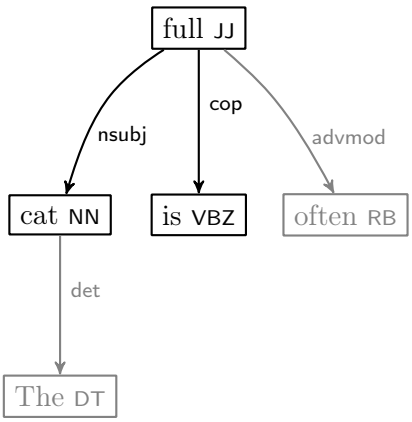

Input text The cat is often full.

Query

```
(fresh [subject adj cop]
  (depends subject "nsubj" adj)
  (depends cop "cop" adj)
  (≡ triple [subject cop adj]))
```

Triples found

cat	be	full

This triple builder captures adjectives of subjects, e.g. that the *cat is full*.

We can qualify the object *full* with an additional triple. See the next triple builder nsubj-advmod.

[24] verbs with no object

nsubj-advmod

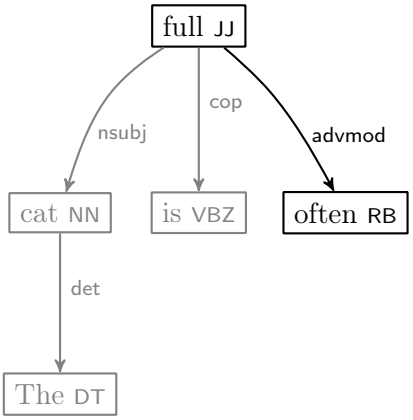

Input text The cat is often full.

Query

```
(fresh [subject predicate advmod cop]
  (depends subject "nsubj" predicate)
  (depends advmod "advmod" predicate)
  (cond^a
   [(verb° predicate)
    (≡ triple [subject predicate advmod])]
   [(≡ triple [predicate :be advmod])]))
```

Triples found

full	:be	often

In combinations with nsubj-adj-cop finds additional descriptions of an adjective, but also additional properties of subjects.
Following e.g. the triple builders nsubj-amod or nsubj-VB, :be is used as our predicate again.
The object from the triple found in the example of nsubj-adj-cop *full* is qualified here with *often*.

nsubj-pred-acomp

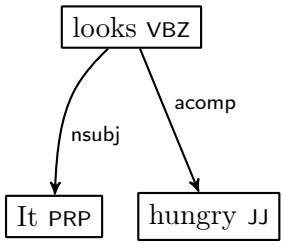

Input text It looks hungry.

Query

```
(fresh [subject activity acomp]
  (depends subject "nsubj" activity)
  (depends acomp "acomp" activity)
  (≡ triple [subject activity acomp]))
```

Triples found

it	look	hungry

Verbs and adjectives have an acomp (adjectival complement) relation, if an adjective can be treated as the verb's object [5].
The triple extraction is straightforward. The adjective becomes the triple's object. The verb's subject the triple's subject and the verb itself the predicate.

nsubj-pred-xcomp

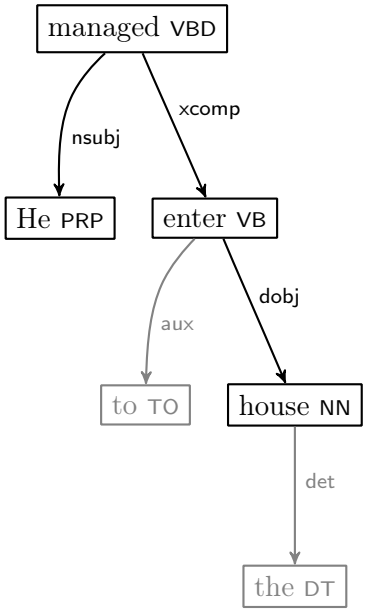

Input text He managed to enter the house.

Query

```
(fresh [subject activity xcomp object relation]
  (depends subject "nsubj" activity)
  (depends xcomp "xcomp" activity)
  (depends object relation xcomp)
  (member° relation ["advmod" "dobj"])
  (cond^e
   [(≡ triple [subject activity xcomp])]
   [(≡ triple [subject xcomp object])]))
```

Triples found

he	manage	enter
he	enter	house

enter is an open clausal complement (xcomp) of *managed*. *enter* does not have its own subject but refers to the subject of *managed* (*He*) [5].
There are two useful triples that can be extracted from the graph. One is the fact that *He managed* to do something (*enter*) and one is a fact about what *He* entered (the *house*).

nsubjpass-pred-agent

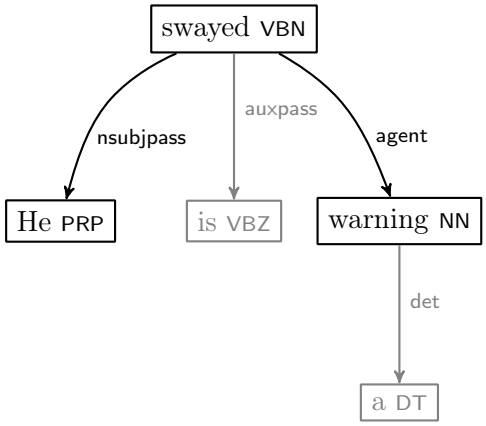

Input text He is swayed by a warning.

Query

```
(fresh [subject predicate object]
  (depends subject "nsubjpass" predicate)
  (depends object "agent" predicate)
  (≡ triple [object predicate subject]))
```

Triples found

warning	sway	he

An agent is "introduced by the preposition *by*" [5]. This also implies a passive subject (*he*), which we will use for our triple's object. The agent (*warning*) is used as the triple's subject, because it does something with the passive subject.

agent-ccomp-dobj

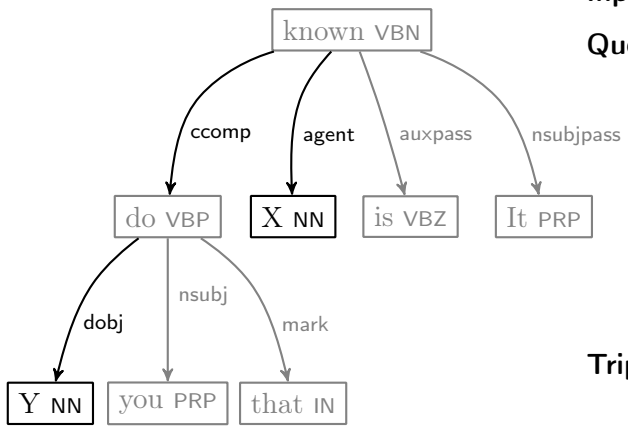

Input text It is known by X that you do Y

Query

```
(fresh [agent object predicate ccomp]
  (depends agent "agent" predicate)
  (depends ccomp "ccomp" predicate)
  (depends object "dobj" ccomp)
  (≡ triple [agent :about object]))
```

Triples found

x	:about	y

This triple builder extracts the fact that *X* talks about *Y*. Because there is no direct graph node we could use for the predicate here, we introduce another derived predicate :about.
Any other information from this sentence, will be captured by other triple builders.

possessive

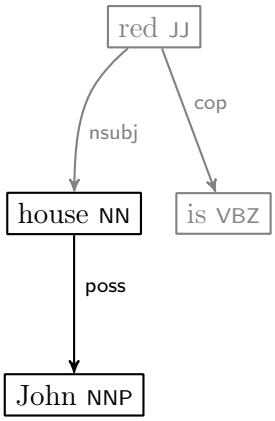

Input text John's house is red.

Query

```
(fresh [subject object]
  (depends subject "poss" object)
  (noun° subject)
  (noun° object)
  (≡ triple [subject :have object]))
```

Triples found

John	:have	house

We want to determine what kind of possessions a subject has. The semantic graph has the relation *poss* (possession modifier) for this. Because we are missing a direct predicate in the graph, we introduce a derived predicate :have.

nsubjpass-ccomp

Input text X and Y are connected as are B and C

Query

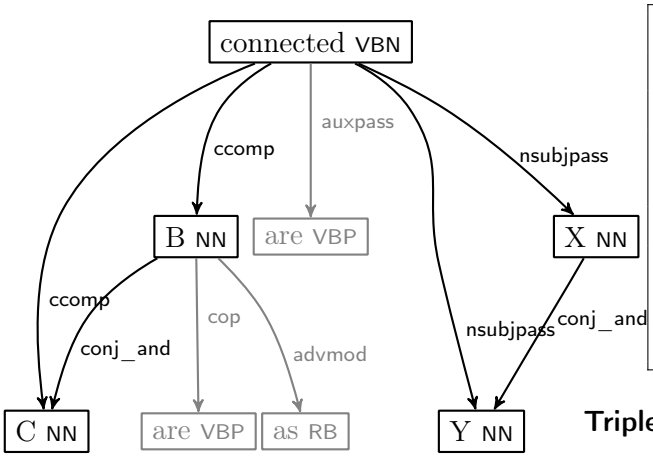


```
(fresh [subj1 predicate subj2 reln1 reln2]
  (depends subj1 reln1 predicate)
  (depends subj2 reln2 predicate)
  (member° reln1 ["nsubjpass" "ccomp"])
  (member° reln2 ["nsubjpass" "ccomp"])
  (depends subj1 ["conj" "and"] subj2)
  (cond^e
    [(≡ triple [subj1 predicate subj2])]
    [(≡ triple [subj2 predicate subj1])]))
```

Triples found

c	connect	b
b	connect	c
y	connect	x
x	connect	y

The query searches for nouns that are connected via some predicate and have a clausal complement (*ccomp*) or passive nominal subject (*nsubjpass*) relation with it.
Because the subject/object-order of the nouns in the resulting triple shouldn't matter, the query is allowed to succeed for all possible combinations of them, with the additional constraint that the nouns have to be connected via some conjunction. This prevents extracting wrong triples, like e.g. "b connect x".

Counterexample. This triple builder sometimes finds information that is obviously wrong. If we vary the sentence a little bit by changing "connected" to "proven", we get essentially the same semantic graph and triples with wrong information.
Running the triple builder on the sentence „X and Y are proven as are B and C" would return the triples

y	prove	x
x	prove	y
c	prove	b
b	prove	c

Clearly this is wrong, "proven" is an intransitive verb here and X and Y did not "prove" each other, but were proven by some (in this case unknown) agent.
There is something missing here. Integrating a verb ontology might help to determine if a verb is transitive or intransitive.
In the case of intransitive verbs a more appropriate result (matching the triples found by the triple builder nsubj-VB) would be:

x	:be	prove
y	:be	prove
c	:be	prove
b	:be	prove

prep-noun

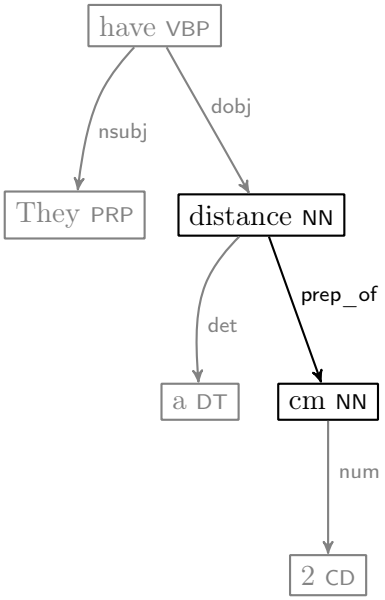

Input text They have a distance of 2 cm.

Query

```
(fresh [obj subj activity prep]
  (depends obj ["prep" prep] subj)
  (noun° obj)
  (noun° subj)
  (project [prep]
    (≡ triple [subj (keyword prep) obj])))
```

Triples found

distance	:of	cm

This triple builder captures preposition between two nouns (or pronouns). Because CoreNLP collapses prepositions, we introduce derived predicates for each preposition. In this case because of the edge between the subject and object (prep_of) the predicate will be :of.

noun-prep-noun

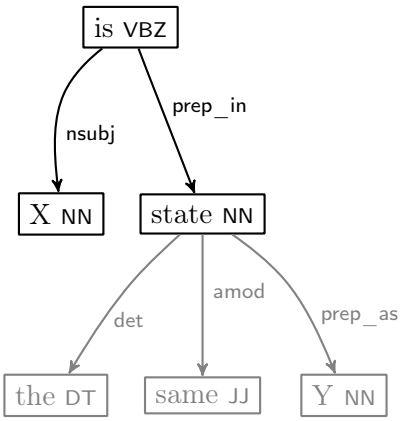

Input text X is in the same state as Y

Query

```
(fresh [obj subj prep activity]
  (depends subj "nsubj" activity)
  (depends obj ["prep" prep] activity)
  (noun° obj)
  (noun° subj)
  (verb° activity)
  (cond^e
   [(project [prep]
      (≡ triple [subj (keyword prep) obj]))]
   [(≡ triple [subj activity obj])]))
```

Triples found

x	be	state
x	:in	state

The complement to prep-noun that captures preposition that are indirectly connected to a noun. Like in prep-noun we again use a derived predicate for capturing the preposition. Because the nouns are not connected directly, we let the query succeed twice. Once for capturing the preposition as our predicate (:in) and once for capturing the actual predicate from the graph (*is*).

noun-num

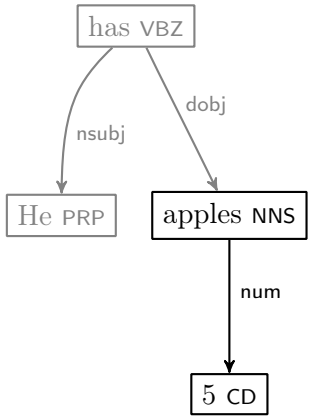

Input text He has 5 apples.

Query

```
(fresh [unit num]
  (depends num "num" unit)
  (noun° unit)
  (tag° num ["CD"])
  (≡ triple [unit :be num]))
```

Triples found

apple	:be	5

This triple builder captures numeric modifiers (num) [5] of nouns, e.g. how many of *apple*s there are (*5*).
Because we are missing a direct predicate in the graph, we will use the derived predicate :be.

advmod-npadvmod-num

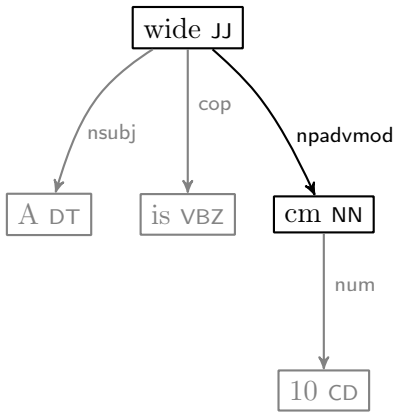

Input text A is 10 cm wide.

Query

```
(fresh [advmod unit num]
  (cond^e
   [(depends advmod "advmod" (lvar))]
   [(word-map advmod)
    (featurec advmod {:tag  "JJ"})])
  (depends unit "npadvmod" advmod)
  (depends num "num" unit)
  (≡ triple [advmod :be unit]))
```

Triples found

wide	:be	cm

The complement to noun-num to capture the unit of a measurement.

Table 4.1 shows all triples that are captured by running them on an example text that is used in Chapter 5 to draw a graph from the text.

Looking at the table it is clear that all necessary information for this is available, however we currently do not have any way of discerning if an instance of e.g. “distance” is the same as another instance of “distance” or what “it” refers to.

4.2. Reifying triples

Currently the triples' subjects or objects refer to one word map only. For example, looking at Table 4.1 we see several subjects with a "Test" or "it" lemma. There is no way that we can know if the instances refer to the same entity or to distinct entities.

By using the same-as relation we build in Section 3.4, we can group the triples' subjects and objects with all other words in the knowledge base that can be treated as refering to the same entity, reifying the triples.

If we apply same-as to every subject and object of our triples our table might look like Table 4.2, where the subjects and objects are replaced by the word groups that the original word belonged to.

Example. Let us look at the coreference clusters found by CoreNLP for our input text. This gives us an approximation of what the same-as relation looks like:

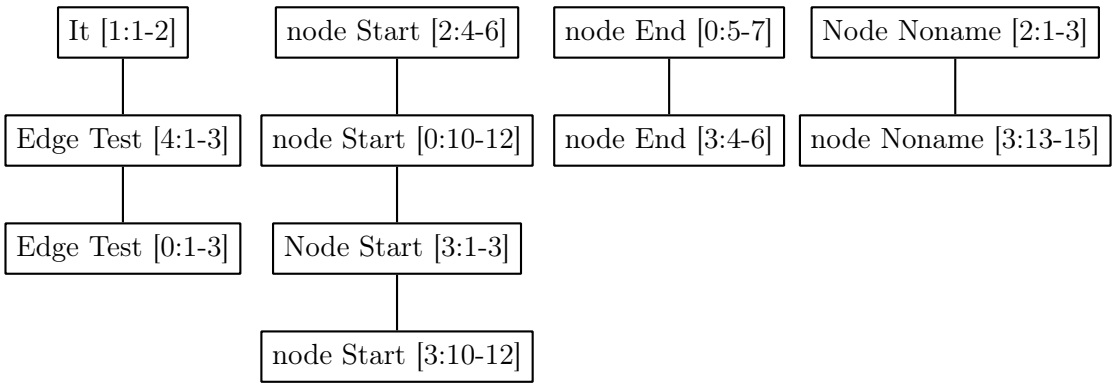

If we look at the subject word map w of the triple

Test go End

And we search for every v for which (same-as w v) holds, we get the word group:[25]

Edge (0,1), Test (0,2), it (1,1), Edge (4,1)

Because working with a list of word maps is cumbersome and also hard to refer to, we give each word group a unique name. A word group's name is made up of the nouns (this excludes *Wh*-words or pronouns) of each word in the group. Because some names may not be unique, we add a number to it. We increment the number every time there is a duplicate name for a group.

A predicate can be derived or refer to a concrete word map. We simply reify the predicates by making them all derived. For non-derived predicates we use it's lemma and derived predicates are copied verbatim. This essentially also groups the predicates together.

The result of applying these steps to Table 4.2 can be seen in Table 4.3.

Example. There are three word groups with the same name: cm-0, cm-1, cm-2
Because they all refer to separate entities, they all have a unique number.

[25] The numbers correspond to the word maps sentence and token index (*sentence*, *index*).

Example. In Table 4.2 the predicate *connect* occurs several times as (potentially distinct) word maps. In Table 4.3 we reduced them all to a single :connect.

To allow clients to access the reified triple, we update the knowledge base from Section 3.4 with a new relation:

(triple t)

triple is a unary relation and we simply add every triple we found to it. Clients can then access any information from the triple by unifying with t. A reified triple is a map that has the following schema:

{:subject {:symbol *word group name*
:group *vector of word maps*}
:predicate *predicate*
:object *same layout as :subject*}

Table 4.1.: Result of running all triple builders on the input text:
`Edge Test goes to node End and starts at node Start. It is 15 cm long. Node Noname and node Start have a distance of 10 cm. Node Start and node End are connected as are node Start and node Noname. Edge Test is 4 cm long.`
Each row shows a captured triple with its subject, predicate and object and the triple builder that created it.

Subject	**Predicate**	**Object**	**Triple builder**
Start	have	distance	:nsubj-pred-dobj
Noname	have	distance	:nsubj-pred-dobj
Start	have	distance	:nsubj-pred-dobj
distance	:of	cm	:prep-noun
long	:be	cm	:advmod-npadvmod-num
cm	:be	4	:noun-num
Test	go	End	:noun-prep-noun
Test	:to	End	:noun-prep-noun
long	:be	cm	:advmod-npadvmod-num
cm	:be	10	:noun-num
Test	start	Start	:noun-prep-noun
Test	:at	Start	:noun-prep-noun
cm	:be	15	:noun-num
Test	be	long	:nsubj-advmod
it	be	long	:nsubj-advmod
Noname	connect	Start	:nsubjpass-ccomp
Start	connect	Noname	:nsubjpass-ccomp
End	connect	Start	:nsubjpass-ccomp
Start	connect	End	:nsubjpass-ccomp
Test	:be	be	:nsubj-VB
Noname	:be	have	:nsubj-VB
it	:be	be	:nsubj-VB
Test	:be	start	:nsubj-VB
Start	:be	have	:nsubj-VB
Test	:be	go	:nsubj-VB

Table 4.2.: Grouped triples based on Table 4.1

Subject word group	Predicate	Object word group
node (0,10), Start (0,11), node (2,4), Start (2,5), Node (3,1), Start (3,2), node (3,10), Start (3,11)	have	distance (2,8)
Node (2,1), Noname (2,2), node (3,13), Noname (3,14)	have	distance (2,8)
node (0,10), Start (0,11), node (2,4), Start (2,5), Node (3,1), Start (3,2), node (3,10), Start (3,11)	have	distance (2,8)
distance (2,8)	:of	cm (2,11)
long (4,6)	:be	cm (4,5)
cm (4,5)	:be	4 (4,4)
Edge (0,1), Test (0,2), it (1,1), Edge (4,1), Test (4,2)	go	node (0,5), end (0,6), node (3,4), end (3,5)
Edge (0,1), Test (0,2), it (1,1), Edge (4,1), Test (4,2)	:to	node (0,5), end (0,6), node (3,4), end (3,5)
long (1,5)	:be	cm (1,4)
cm (2,11)	:be	10 (2,10)
Edge (0,1), Test (0,2), it (1,1), Edge (4,1), Test (4,2)	start	node (0,10), Start (0,11), node (2,4), Start (2,5), Node (3,1), Start (3,2), node (3,10), Start (3,11)
Edge (0,1), Test (0,2), it (1,1), Edge (4,1), Test (4,2)	:at	node (0,10), Start (0,11), node (2,4), Start (2,5), Node (3,1), Start (3,2), node (3,10), Start (3,11)
cm (1,4)	:be	15 (1,3)
Edge (0,1), Test (0,2), it (1,1), Edge (4,1), Test (4,2)	be	long (4,6)
Edge (0,1), Test (0,2), it (1,1), Edge (4,1), Test (4,2)	be	long (1,5)
Node (2,1), Noname (2,2), node (3,13), Noname (3,14)	connect	node (0,10), Start (0,11), node (2,4), Start (2,5), Node (3,1), Start (3,2), node (3,10), Start (3,11)
node (0,10), Start (0,11), node (2,4), Start (2,5), Node (3,1), Start (3,2), node (3,10), Start (3,11)	connect	Node (2,1), Noname (2,2), node (3,13), Noname (3,14)
node (0,5), end (0,6), node (3,4), end (3,5)	connect	node (0,10), Start (0,11), node (2,4), Start (2,5), Node (3,1), Start (3,2), node (3,10), Start (3,11)
node (0,10), Start (0,11), node (2,4), Start (2,5), Node (3,1), Start (3,2), node (3,10), Start (3,11)	connect	node (0,5), end (0,6), node (3,4), end (3,5)
Edge (0,1), Test (0,2), it (1,1), Edge (4,1), Test (4,2)	:be	be (4,3)
Node (2,1), Noname (2,2), node (3,13), Noname (3,14)	:be	have (2,6)
Edge (0,1), Test (0,2), it (1,1), Edge (4,1), Test (4,2)	:be	be (1,2)
Edge (0,1), Test (0,2), it (1,1), Edge (4,1), Test (4,2)	:be	start (0,8)
node (0,10), Start (0,11), node (2,4), Start (2,5), Node (3,1), Start (3,2), node (3,10), Start (3,11)	:be	have (2,6)
Edge (0,1), Test (0,2), it (1,1), Edge (4,1), Test (4,2)	:be	go (0,3)

Table 4.3.: Reified triples based on Table 4.2

Subject	**Predicate**	**Object**
node-start-0	:have	distance-0
node-noname-0	:have	distance-0
distance-0	:of	cm-2
long-0	:be	cm-1
cm-1	:be	num-4-0
edge-test-0	:go	end-node-0
edge-test-0	:to	end-node-0
long-1	:be	cm-0
cm-2	:be	num-10-0
edge-test-0	:start	node-start-0
edge-test-0	:at	node-start-0
cm-0	:be	num-15-0
edge-test-0	:be	long-0
edge-test-0	:be	long-1
node-noname-0	:connect	node-start-0
node-start-0	:connect	node-noname-0
end-node-0	:connect	node-start-0
node-start-0	:connect	end-node-0
edge-test-0	:be	be-1
node-noname-0	:be	have-0
edge-test-0	:be	be-0
edge-test-0	:be	start-0
node-start-0	:be	have-0
edge-test-0	:be	go-0

4.3. Exporting an OWL ontology

We create an OWL class for every subject and object word group of a triple. OWL 2 introduces a feature called *punning* [9] where we create an individual for each of our classes and then use object properties on these individuals to describe relationships between the classes.

Because a triple's predicate describes a relationship between its subject and object, we use an object property with a name based on the predicate and add an axiom to the ontology that links the subject's individual with the object's individual via this object property.

A subject or object word group contains word maps that have additional information about that group, such as all of the actual words (tokens) that make up the group. We add this information as datatype properties to the subject's or object's individual.

Figure 4.1 shows a subset of an ontology that is based on Table 4.3. The "has individual" loops are an artifact of our use of punning. Even though node-start-0, node-noname-0 and num-10-0 are only shown as individuals, they are still represented as classes (and as subclasses of Thing) in the full ontology.

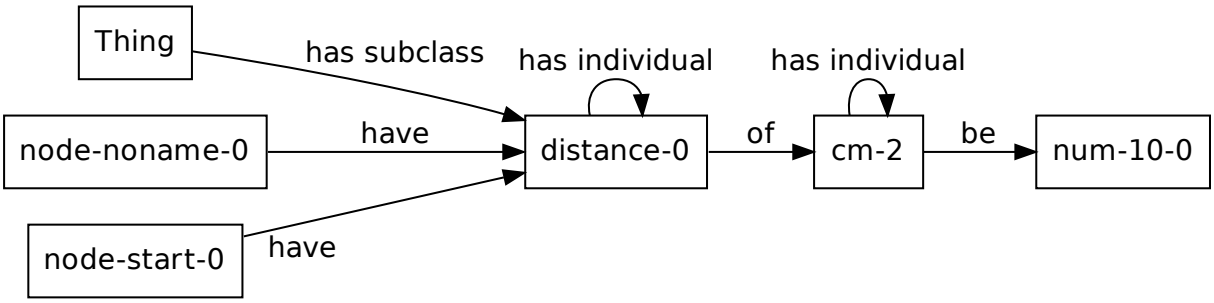

Figure 4.1.: A view on the ontology that is extracted from Table 4.3

Example. We can query the ontology using SparQL. E.g. to return all subject word groups with their constituent tokens that are linked to the word group distance-0, we can run the following query:

```
1 PREFIX : <http://clide.informatik.uni-bremen.de/clide-nlp#>
2 SELECT ?subject ?token WHERE {
3  ?subject :have :distance-0.
4  ?subject :hasToken ?token
5 }
```

Query result:

?subject	?token
node-noname-0	"Noname"
node-noname-0	"Node"
node-noname-0	"node"
node-start-0	"Start"
node-start-0	"node"

5. Use case: Graph creation from a natural language specification

This chapter introduces an example application that can create graphs from a text. The application is directly integrated into *clide-nlp* and uses the reified triples introduced in Section 4.2 as its input.

The application and its input can be seen in the reconciler's dependency graph Figure 3.2 and is named :draw there.

The output is a graph with the nodes and edges as described in the text and a list of warnings about ambiguous or incomplete information extracted from the triples.

The application should detect

- named nodes,
- edges that are directly specified with a name,
- edges that are indirectly specified as a connection between two nodes, and
- edge lengths.

Emit warnings when there are

- simple contradictions, like different lengths for the same edge,
- unfinished edge or node specifications,
- sentences that specify the same node or edge twice, and
- non-integer edge lengths.

To achieve these goals, we must specify what kind of sentences we would like to understand. We can specify edges with the following sentences:

Type	**Sentence example**
Unnamed edges	Node A and Node B are connected.
Named edges	Edge B starts at Node A and goes to Node B.
Edge distance	Edge B is 5 cm long.
Edge distance	Node A and Node B have a distance of 5 cm.

Nodes are specified implicitly by mentioned e.g. "Node A" somewhere in the text.

Distances have a unit and magnitude. We only support "cm" as a unit.

Because we use triples as the basis of our application, we can structure the sentences differently,

while keeping the triple set the same.

Example. Extracting triples from the sentences

> Node A and Node B are connected. Node A and Node B have a distance of 5 cm.

and the sentence

> Node A and Node B are connected with a distance of 5 cm.

will result in the same set of triples.

5.1. Triple walks

We define some helper goals that let us define a walk that follows a chain of triples, and allows us to essentially pattern match on that chain.

(subject$^o\rightarrow$ t & *clauses*)

A subject walk succeeds if $t = (S_0, P_0, O_0)$ can satisfy each of the *clauses*. t is the triple we start with.

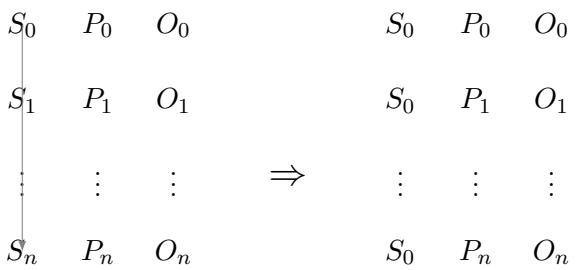

Figure 5.1.: Illustration of how a subject walk beginning with triple t, threads t's subject through the whole triple chain.

A clause can be a tuple with a predicate and object or a 3-tuple in which case the last element is a partial match of a word map in the S_0's word group.
We *thread* the subject of t through each of the clauses. This is illustrated in Figure 5.1.

Example. We define the following subject walk to find an edge:

S_0	:start	O_1	where a word map in S_0 must match {:lemma "Edge"}
S_0	:at	O_1	
S_0	:go	O_2	
S_0	:to	O_2	

We use O_1 and O_2 in multiple clauses to make sure that the clauses only match the same object group.

Looking at Table 4.3 we can find a triple for which the walk succeeds: t = (edge-test-0, :start, node-start-0)

edge-test-0	:start	node-start-0
edge-test-0	:at	node-start-0
edge-test-0	:go	end-node-0
edge-test-0	:to	end-node-0

We can run that walk with subject$^o\rightarrow$ inside a core.logic query:

```
1 (subject^o→ t
2             [:start O_1 {:lemma "Edge"}]
3             [:at O_1]
4             [:go O_2]
5             [:to O_2])
```

If the subject$^o\rightarrow$ goal succeeds, O_1 will be bound to node-start-node-0 and O_2 to end-node-0. We can then extract more information out of them, if necessary and perform some additional validation.

(object$^o\rightarrow$ t & *clauses*)

The counterpart to subject$^o\rightarrow$ that matches on the triples' objects first (see Figure 5.2).

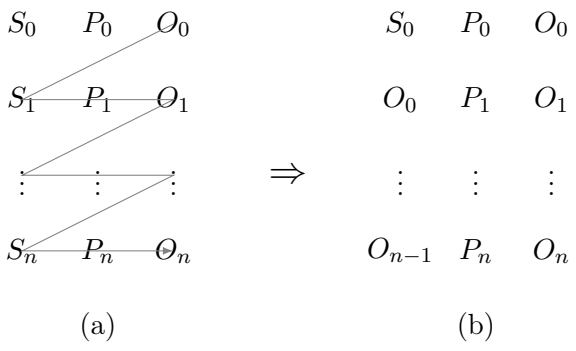

Figure 5.2.: Illustration of how an object walk beginning with triple t, threads the object of each consecutive triple through the triple chain.

Example. We define an object walk to get the distance between the nodes of our edge. We assume that the previous subject walk succeeded with $t = (S_0, P_0, O_0)$.

O_0	:have	O_3	where a word map in O_3 must match {:lemma "distance"}
O_3	:of	O_4	
O_4	:be	O_5	where a word map in O_5 must match {:tag "CD"}

The walk succeeds, if we start with the triple $t = (\text{edge-test-0}, \text{:start}, \text{node-start-0})$ and follow the triples:

node-start-0	:have	distance-0
distance-0	:of	cm-2
cm-2	:be	num-10-0

We can run that walk with object$^o\!\rightarrow$ inside a core.logic query:

```
1 (object^o→ t
2             [:have  O_3 {:lemma  "distance"}]
3             [:of O_4]
4             [:be  O_5 {:tag  "CD"}])
```

If the object$^o\!\rightarrow$ goal succeeds and O_3, O_4, and O_5 are logic variables, they will be bound to distance-0, cm-2, and num-10-0 respectively.
We can then extract more information out of them, if necessary.

5.2. Implementation

We implement several collection stages that search the triples for relevant information. The general process follows these steps:

1. Collect node
2. Collect edges
3. Check found edges for inconsistencies
4. Check for singleton nodes

Each stage might emit warnings about inconsistencies, which we need to present to the user later.

The node and edge collecting stages make use of subject and object walks to extract the information that we need to create a graph.

Example. If we combine the object walk with the subject walk in our previous examples, we get an edge with the distance between its nodes:

- The group $S_0 =$ edge-test-0 contains the edge label (*Test*)
- The group $O_5 =$ num-10-0 is the edge's length
- The group $O_4 =$ cm-2 contains the edge length's unit

An edge's or node's label is extracted by using the word that is immediately next to it in the original text.
Extracting the edge's length and unit is easy, because O_5 and O_4 are word groups with only a single word map, so we can simply use that word map's lemma. We then check if the length is an integer and the unit is one of the supported units (cm).

The output of our example application is presented to the users as additional annotations:

draw — Level: Chunk

```
1 Edge Test goes to node End and starts at node Start.
2 It is 2 cm long.
3 Node Noname and node Start have a distance of 1 cm.
4 Node Start and node End are connected as are node Start and
5 node Noname. Edge Test is 1 cm long.
6
```

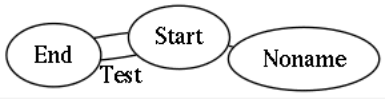

Shows the graph that was extracted from the chunk's text. The input text contains some ambiguities, which means the graph is not looking like we want it to look. The double edge between nodes "Start" and "End" looks especially suspicious.

draw-warnings and draw-warning-highlights — Level: Chunk

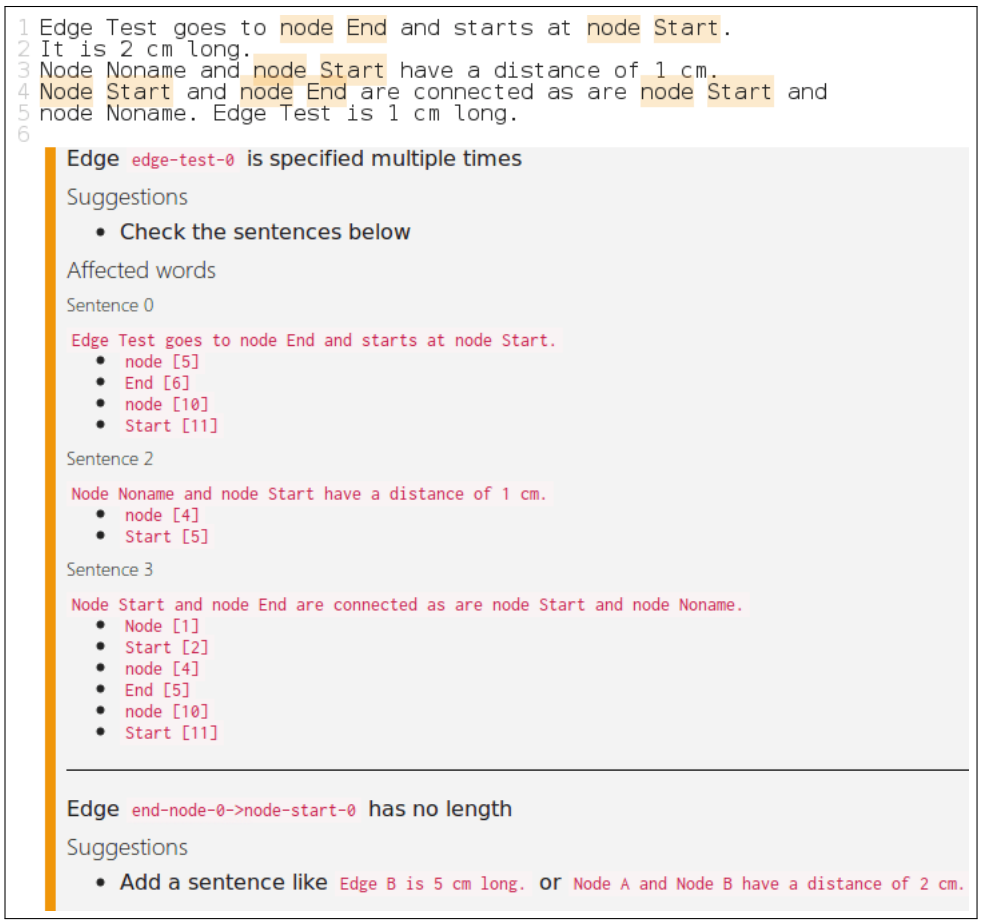
1 Edge Test goes to node End and starts at node Start.
2 It is 2 cm long.
3 Node Noname and node Start have a distance of 1 cm.
4 Node Start and node End are connected as are node Start and
5 node Noname. Edge Test is 1 cm long.
6
Edge edge-test-0 is specified multiple times
Suggestions
Check the sentences below
Affected words
Sentence 0
Edge Test goes to node End and starts at node Start.
node [5]
End [6]
node [10]
Start [11]
Sentence 2
Node Noname and node Start have a distance of 1 cm.
node [4]
Start [5]
Sentence 3
Node Start and node End are connected as are node Start and node Noname.
Node [1]
Start [2]
node [4]
End [5]
node [10]
Start [11]
Edge end-node-0->node-start-0 has no length
Suggestions
Add a sentence like Edge B is 5 cm long. or Node A and Node B have a distance of 2 cm.

The information collecting stages emit warnings about potential problem areas in our input text. We display some information about which sentences and words might be problematic and some suggestions about how to resolve the warnings.

Here we can see that the edge “Test” is specified twice which explains the double edge we saw earlier.

6. Conclusion

We have shown how we can extract information from a text in a straightforward way, while only using the simple tools provided by CoreNLP, and what kind of problems result in integrating such a system into a distributed development environment like *Clide*.

To achieve our initial goals, we have taken the following approach:

1. We provide an underlying NLP knowledge base (see Section 3.3) for an input text based on CoreNLP's dependency parser and its coreference resolution system.
2. The reconciler (see Section 3.2) makes sure that text changes sent by *Clide* are integrated into the knowledge base, and keeps the ontology and annotations up to date and in sync with the input text while making sure to only update them when really necessary. Because coreference resolution is a slow process, and we need to rerun it after each text change, we split the input text into chunks and only operate on one chunk at a time.
3. Triple builders (see Chapter 4) are simple queries on the knowledge base that extract meaningful units of information from the text's semantic graphs in the form of triples $(subject, predicate, object)$.
4. The triples extracted from the queries are augmented by grouping their subjects and objects according to the coreference chain they belong to, yielding an ontology with unique classes (word groups). We provide an example for how to export an OWL ontology (see Section 4.3), so that the ontology can be used by other tools.
5. We expose the underlying structure of sentences and texts to the user by providing annotations that visualize that structure in *Clide* (see Section 3.5).
6. In Chapter 5 we have shown how the ontology can be used to create graphs from a simple natural language specification and how using triples enables sentences in the input text with slightly different phrasing to still yield the same ontology and graph.

There are some (solvable) caveats to our approach:

- Some triples we extract make no sense. This is a symptom of missing information and of the triple builders' simplistic nature. E.g. some triple builders like nsubjpass-ccomp are limited, because there is currently no way to determine a verb's transitivity.
- At the moment we ignore all tenses and merge potentially different states of word groups in different time frames (as indicated e.g. by a triple's predicate's tense) into one ontology. Using a verb ontology in combination with the predicates' part-of-speech tags, we could split the ontology into separate ontologies, one for each time frame. We could then e.g.

track the changes of a word group's attributes over time.

- Currently the ontology we create has no concept hierarchy. We could integrate WordNet to group word groups that have the same concept behind them under one umbrella via e.g. *is-a* relationships in the ontology.
- At the moment the ontologies we extract are limited to one chunk of a text only. It should be possible to develop some heuristics that would allow us to merge the ontologies of two or more chunks together into one ontology.
- *clide-nlp* is limited by *Clide*'s current behavior to keep annotations static and to limit the assistants direct influence to the server-side only. A client side integration could enable some interactive aspects, like e.g. defining a new triple builder on the fly and using Clojure's dynamic aspects to make it available to the system immediately.[26]

The tools we created while building the assistant are general enough to be used outside of *Clide* and could be integrated into other editing environments. Our NLP knowledge base with information from CoreNLP's semantic graphs, its coreference resolution system, and the ontology we extracted from them, provides easy access to information about a text.

Exposing Clojure's and *clide-nlp*'s dynamic natures in *Clide*'s interface and combining it with *Clide*'s collaborative aspects can enable an environment and framework where we can quickly and collaboratively develop simple systems that have some albeit limited and domain specific text understanding.

[26] This is possible already, but not exposed to the user in the UI.

A. Part of Speech Tags

The part of speech tags used by CoreNLP are based on the tags used by the Penn Treebank [15]. There are however some differences. Table A.1 shows an updated version of the part of speech tag table in [15, p. 317] based on experience with CoreNLP. As such this table is most likely incomplete, but enough to follow the examples in this report.

Table A.1.: An incomplete list of part of speech tags used by CoreNLP based on [15, p. 317]

CC	Coordinating conjunction	TO	*to*
CD	Cardinal number	UH	interjection
DT	Determiner	VB	Verb, base form
EX	Existential *there*	VBD	Verb, past tense
FW	Foreign word	VBG	Verb, gerund/present participle
IN	Preposition/subordinating conjunction	VBN	Verb, past participle
JJ	Adjective	VBP	Verb, non-3rd ps. sing. present
JJR	Adjective, comparative	VBZ	Verb, 3rd ps. sing. present
JJS	Adjective, superlative	WDT	*wh*-determiner
LS	List item marker	WP	*wh*-pronoun
MD	Modal	WP$	Possessive *wh*-pronoun
NN	Noun, singular or mass	WRB	*wh*-adverb
NNS	Noun, plural	RP	Participle
NNP	Proper noun, singular	$	Currency sign
NNPS	Proper noun, plural	.	Sentence-final punctuation ! ? .
PDT	Predeterminer	,	Comma
POS	Possessive ending	:	Colon, semi-colon
PRP	Possessive pronoun	-LRB-	Left bracket ([{ character
PP$	Possessive pronoun	-RRB-	Right bracket)] } character
RB	Adverb	``	Left (double or single) quote
RBR	Adverb, comparative	''	Right (double or single) quote
RBS	Adverb, superlative		

B. Installation notes

The CD that accompanies this report contains three ZIP files:

Filename	Contents
`report.pdf`	A digital copy of this report
`clide-nlp-src.zip`	The source code for *clide-nlp*
`clide-src.zip`	The source code to a version of *Clide* that works correctly with *clide-nlp*
`clide-nlp.zip`	Contains an executable JAR of *clide-nlp* and startup scripts

clide-nlp requires Java 7 and *Clide* works best with a WebKit-based browser like e.g. Chrome. You need approx. 3 GiB of RAM to successfully run *clide-nlp*. To execute it:

- Extract `clide-nlp.zip` and run `run.sh` on Linux/FreeBSD or `run.bat` on Windows.
- Wait a minute or two.
- A launcher window will pop up that informs you about the startup process.
- After the system is ready, *clide-nlp* should inform you that it is ready at `http://localhost:14000`
- Open the URL and log in with user `clide-nlp` and password `clide-nlp`.
- Open the `clide-nlp/Example` project and look at `00-README.txt` for further help.
- And most importantly, try editing one of the files!

List of Figures

1.1 The *clide-nlp* annotation semantic-graph showing the semantic graph for the current sentence (highlighted in blue). 9
1.2 The *clide-nlp* annotation reified-triples showing the triples that are found for the given text and the annotation same-as highlighting all words that are detected as belonging together in red. The highlighted words are part of the group Edge-Test-it-0. 10
1.3 Showing all annotation from the example application draw, draw-warnings and draw-warning-highlights that draws the graph specified in the text. 11

3.1 Architecture overview 18
3.2 Chunk annotation dependency graph 22

4.1 A view on the ontology that is extracted from Table 4.3 49

5.1 Illustration of how a subject walk beginning with triple t, threads t's subject through the whole triple chain. 52
5.2 Illustration of how an object walk beginning with triple t, threads the object of each consecutive triple through the triple chain. 53

Bibliography

[1] Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., and Ives, Z. DBpedia: A Nucleus for a Web of Open Data. In *The Semantic Web*. Springer, 2007, pp. 722–735.

[2] Byrd, W. E. *Relational programming in miniKanren: Techniques, applications, and implementations*. PhD thesis, Indiana University, 2010.

[3] Chklovski, T., and Pantel, P. VerbOcean: Mining the Web for Fine-Grained Semantic Verb Relations. In *EMNLP* (2004), vol. 2004, pp. 33–40.

[4] De Marneffe, M.-C., MacCartney, B., Manning, C. D., et al. Generating typed dependency parses from phrase structure parses. In *Proceedings of LREC* (2006), vol. 6, pp. 449–454.

[5] de Marneffe, M.-C., and Manning, C. D. Stanford typed dependencies manual. Manual, Stanford NLP Group, Stanford University, September 2008. Revised in December 2013 for the Stanford Parser v. 3.3.

[6] Emerick, C., Carper, B., and Grand, C. *Clojure Programming*. O'Reilly Media, Inc., 2012.

[7] Fogus, M., and Houser, C. *The Joy of Clojure: Thinking the Clojure Way*. Manning Publications Co., 2011.

[8] Friedman, D. P., Byrd, W. E., and Kiselyov, O. *The Reasoned Schemer*. MIT Press, 2005.

[9] Golbreich, C., Wallace, E. K., and Patel-Schneider, P. F. OWL 2 Web Ontology Language: New Features and Rationale. *W3C Recommendation* (2012).

[10] Hitzler, P. *Semantic Web: Grundlagen*. eXamen.press. Springer-Verlag, Berlin, Heidelberg, 2008.

[11] Lee, H., Chang, A., Peirsman, Y., Chambers, N., Surdeanu, M., and Jurafsky, D. Deterministic coreference resolution based on entity-centric, precision-ranked rules.

[12] Lee, H., Peirsman, Y., Chang, A., Chambers, N., Surdeanu, M., and Jurafsky, D. Stanford's multi-pass sieve coreference resolution system at the CoNLL-2011 shared task. In *Proceedings of the Fifteenth Conference on Computational Natural Language Learning: Shared Task* (2011), Association for Computational Linguistics, pp. 28–34.

[13] Lord, P. The Semantic Web takes Wing: Programming Ontologies with Tawny-OWL. *arXiv preprint arXiv:1303.0213* (2013).

[14] Lüth, C., and Ring, M. A web interface for Isabelle: The Next Generation. In *Intelligent Computer Mathematics*. Springer, 2013, pp. 326–329.

[15] Marcus, M. P., Marcinkiewicz, M. A., and Santorini, B. Building a large annotated corpus of English: The Penn Treebank. *Computational linguistics 19*, 2 (1993), 313–330.

[16] Miller, G. A. WordNet: A Lexical Database for English. *Communications of the ACM 38*, 11 (1995), 39–41.

[17] Raghunathan, K., Lee, H., Rangarajan, S., Chambers, N., Surdeanu, M., Jurafsky, D., and Manning, C. A multi-pass sieve for coreference resolution. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing* (2010), Association for Computational Linguistics, pp. 492–501.

[18] Recasens, M., de Marneffe, M.-C., and Potts, C. The life and death of discourse entities: Identifying singleton mentions. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (2013), pp. 627–633.

[19] Ring, M. Eine webbasierte Entwicklungsumgebung für den interaktiven Theorembeweiser Isabelle. Diplomarbeit, Universität Bremen, 2013.

[20] Ring, M., and Lüth, C. Collaborative Interactive Theorem Proving with Clide. In *Proceedings of ITP 2014* (2014), Springer.

[21] Varjú, Z., Littauer, R., and Ernis, P. Using Clojure in Linguistic Computing. In *Proceedings of the 5th European Lisp Symposium* (2012).

[22] Wynne, M., and Hellesoy, A. *The Cucumber Book: Behaviour-driven Development for Testers and Developers*. Pragmatic Bookshelf, 2012.